

LAMP-TR-091
CS-TR-4391
UMIACS-TR-2002-71

August 2002

Constraints on the Generation of Tense, Aspect, and Connecting Words from Temporal Expressions

Bonnie Dorr, Terry Gaasterland

Language and Media Processing Laboratory
Institute for Advanced Computer Studies
College Park, MD 20742

Abstract

Generating language that reflects the temporal organization of represented knowledge requires a language generation model that integrates contemporary theories of tense and aspect, temporal representations, and methods to plan text. This paper presents a model that produces event combinations and appropriate connecting words to relate them. We distinguish between inherent and non-inherent aspectual features of verbs and describe an algorithm that uses these features to select tense, aspect, and temporal connecting words for generating text based on time-stamped information. The main result of this work is the successful incorporation of constrained linguistic theories of tense and aspect in a self-contained module called CONGEN that produces a ranked list of temporal connectives and tense/aspect possibilities from pairs of time-stamped literals. We show that the theoretical results described herein have been verified in a large-scale corpus analysis. The framework serves as the basis of a component designed to enhance the English output of a constrained generation system.

***The support of the LAMP Technical Report Series and the partial support of this research by the National Science Foundation under grant EIA0130422 and the Department of Defense under contract MDA9049-C6-1250 is gratefully acknowledged.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE AUG 2002		2. REPORT TYPE		3. DATES COVERED 00-08-2002 to 00-08-2002	
4. TITLE AND SUBTITLE Constraints on the Generation of Tense, Aspect, and Connecting Words from Temporal Expressions			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Language and Media Processing Laboratory, Institute for Advanced Computer Studies, University of Maryland, College Park, MD, 20742-3275			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 55	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Constraints on the Generation of Tense, Aspect, and Connecting Words from Temporal Expressions

(Submitted to JAIR, August, 2002)

Bonnie J. Dorr¹ and Terry Gaasterland²

¹ Department of Computer Science
A.V. Williams Building
University of Maryland
College Park, MD, 20742
bonnie@umiacs.umd.edu
<http://umiacs.umd.edu/labs/CLIP>

² The Rockefeller University
1230 York Avenue
New York, NY, 10021
gaaster@genomes.rockefeller.edu

Abstract. Generating language that reflects the temporal organization of represented knowledge requires a language generation model that integrates contemporary theories of tense and aspect, temporal representations, and methods to plan text. This paper presents a model that produces event combinations and appropriate connecting words to relate them. We distinguish between inherent and non-inherent aspectual features of verbs and describe an algorithm that uses these features to select tense, aspect, and temporal connecting words for generating text based on time-stamped information. The main result of this work is the successful incorporation of constrained linguistic theories of tense and aspect in a self-contained module called CONGEN that produces a ranked list of temporal connectives and tense/aspect possibilities from pairs of time-stamped literals. We show that the theoretical results described herein have been verified in a large-scale corpus analysis. The framework serves as the basis of a component designed to enhance the English output of a constrained generation system.

Keywords: natural language generation, tense, aspect, connecting words, temporal knowledge

1 Introduction

Reasoning about temporal knowledge for machine translation (MT) and formulating answers to questions (Q/A) about time necessitate the presentation of temporal information. One approach to presenting this information is through

natural language sentences that contain acceptable combinations of temporal expressions. This requires a method to plan language and to make appropriate selections in the surface-level sentence for tense (e.g., past, present, and future), aspect (e.g., perfect and progressive), and temporal connecting word (e.g., *after*, *before*, and *while*). This paper presents a model that incorporates contemporary theories of tense and aspect and develops a new framework for selecting tense, aspect, and temporal connecting words. We explore the interrelationships between choices in each of these categories, and then show how individual selection models—one for aspect, one for tense, and one for connecting words—combine into a single overall approach.

We adopt a constraint-based approach to characterizing temporally encoded input in natural language. Consider the following sentence:

- (1) Mary had caught her plane before John arrived.

This is a natural language characterization of a temporal situation in which Mary’s travel is initiated *before* John’s arrival (sometime in the past).

If John’s arrival has not yet occurred, the above sentence would be ruled out as a possibility because the past tense (i.e., *arrived*) is not temporally consistent with knowledge about John’s arrival (i.e., that it has not yet occurred). Instead, the following natural language characterization might be produced:

- (2) Mary has caught her plane before John arrives.

In addition to constraints based on temporal knowledge, an inappropriate surface realization might be ruled out because of linguistic constraints on tense combinations. For example, the following two cases would be ruled out due to restrictions on the combination of the past tense with the present tense:¹

- (3) (i)* Mary caught her plane before John arrives.
(ii)* Mary had caught her plane before John arrives.

The main result of this work is the successful incorporation of constrained linguistic theories of tense and aspect in a self-contained module called CON-GEN (CONnective GENerator) that produces a ranked list of temporal connectives and tense/aspect possibilities from pairs of time-stamped literals. The module is designed to operate in the larger context of a full text planner or MT system, where the input to the module is formalized as a conjunction of two time-stamped literals and their corresponding verb tokens. Our approach extends earlier work by Dorr and Gaasterland (1995) in that it accommodates tense-pair combinations for events with duration.

A time-stamped *literal* is a logical expression of the form $p(x_1, \dots, x_n, \text{start-time}, \text{stop-time})$ where p is a relation name, each x_i is either a variable or a constant, and the time-stamp is expressed in terms of a start time and stop time (which is ∞ for open-ended events). For example, the time-stamped literal *laugh*(*Mary*, 14:01, 14:03) describes an event in which Mary laughs for two minutes, and *draw*(*John*, *circle*, 14:00, 14:10)

¹ The asterisk (*) indicates ungrammaticality.

describes an event in which John draws a circle for 10 minutes. These literals are in a form that is compatible with representations provided in temporal databases such as those defined by (Androutsopoulos, 1996; Snodgrass, 1995; Snodgrass, 1990; Torp, Jensen, and Snodgrass, 2000).

For readability, the input to CONGEN—a full dependency tree—is formalized in terms of this simple time-stamped literal notation. Prior to processing, it is assumed that the input text has been temporally annotated (possibly in a foreign language, as in the work of Baldwin (2002)) and analysis has taken place to produce a time-stamped dependency tree in the format of the PENMAN Sentence Planning Language (SPL) (Kasper, 1989). If the input is in a foreign language, a simple translation procedure (described in (Habash, 2002)) converts source-language words into sets of English words while maintaining the original dependency structure (which, as mentioned above, is formalizable as one or more time-stamped literals). During processing, CONGEN annotates this converted dependency tree further with ranked connective and tense/aspect possibilities. The result is then made available for processing by a generation-heavy component of an MT system (GHMT) (Dorr, Habash, and Traum, 1998; Habash and Dorr, 2002), which produces the English output by means of “constrained overgeneration.”

Within the GHMT framework, a range of possible verbs are selected for the time-stamped literals; these choices are ranked and the final selection is made based on the highest ranking English sentence. A supporting source of information for that ranking is the lexical-choice component of (Dorr and Voss, 1996; Dorr and Olsen, 1996) where verb choices are narrowed down according to the associated aspectual information. Further improvements on the lexical-selection algorithm are reported in (Olsen et al., 2000; Olsen et al., 2001). Figure 1 illustrates the relation between the components, showing the flow in and out of CONGEN.

A major contribution of our work is the successful application of constrained linguistic theories of tense and aspect to the generation of event combinations and the selection of appropriate connecting words that relate them. We distinguish between inherent and non-inherent aspectual features of verbs and describe an algorithm that uses these features to select tense, aspect, and temporal connecting words for generated text based on time-stamped information. We show that the theoretical results described herein have been verified in a large-scale analysis, using the Lancaster-Oslo-Bergen corpus.²

The following section motivates our approach, setting our work in the context of related temporal frameworks currently under investigation by other researchers. Section 3 provides background on linguistic theories of aspect and tense. Section 4 describes our extension of Hornstein’s theory of tense to handle not only point events but also events with duration. Section 5 describes the algorithm for generating text from temporal expressions and provides details behind selecting aspect and connecting words. Section 6 shows the result of running

² We are indebted to three anonymous reviewers who inspired our corpus-based experimentation so that we could verify our theoretical results.

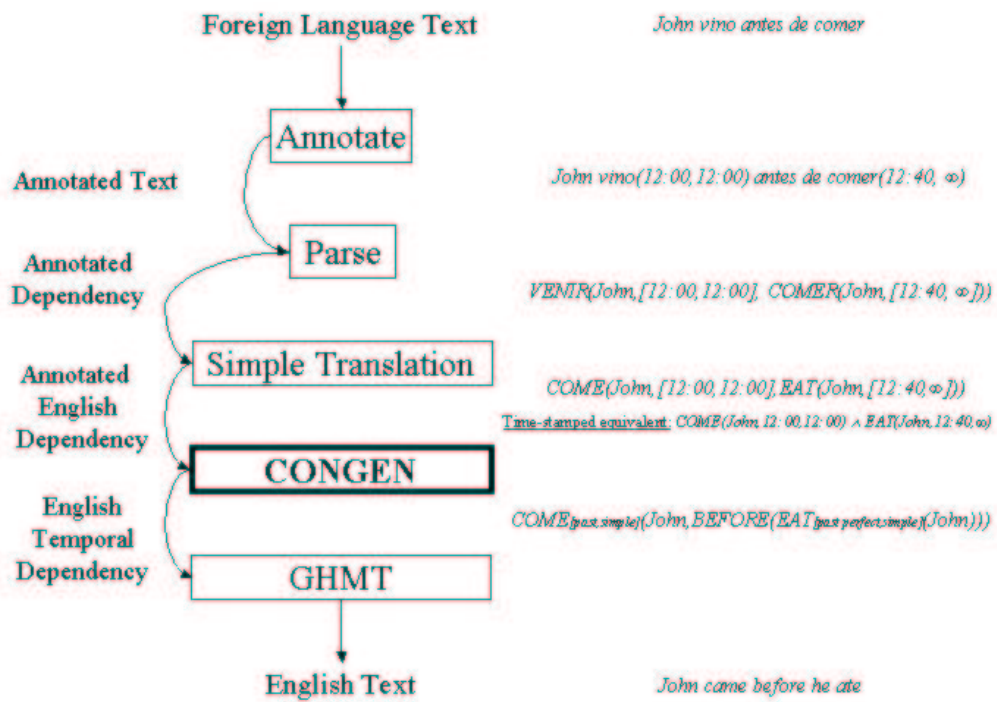


Fig.1. Input and Output of CONGEN in Context of Overall MT System

the algorithm on several specific examples. Section 7 presents a corpus-based analysis used to verify our theoretical results. Finally, Section 8 discusses the implications of this work as well as limitations and future work.

2 Motivation and Related Work

The use of time-stamped literals as our starting point is an approach that is increasingly more critical and relevant, especially with recent developments in generating English sentences, e.g., question-answering from time-stamped data (Pustejovsky, Wiebe, and Maybury, 2002). In the approach of Schilder and Habel (2001) and Filatova and Hovy (2001), time stamps are assigned to temporally related clauses in order to enhance generation. Many researchers are focusing on the development of temporal formalisms such as TimeML for the creation of time-stamped gold standards, including TimeBank (Pustejovsky et al., 2002; Radev and Sundheim, 2002). The work of Setzer (2001), Setzer and Gaizauskas (2001), and Setzer (2002) has resulted in a fine-grained time-stamp formalism for the annotation of events, times, and temporal relations in newswire text. Mani and Wilson (2000) and Wilson et al. (2001) have shown that it is possible to achieve 83.2% accuracy in the annotating of newswire text with time values. More recently, there has been work on language-independent temporal annotation, e.g., the learning system of Baldwin (2002) which induces French annotation rules that are applicable to untagged data.

A common standard for the time stamps used by many of these researchers is the set of TIMEX2 tags, taken originally from TIDES TDT-2 (Ferro et al., 2000). An example is shown here for the sentence *John flew to Japan at 11:30am*:

(4) John flew to Japan
 <TIMEX2 VAL="PRESENT_REF" ANCHOR_VAL="1998-01-30T11:30" ANCHOR_DIR="AS_OF">
 at 11:30am

In our simplified time-stamped literal notation, the above example is rendered as the following representation:

(5) fly(John,Japan,t₁,t₂)

where t₁ is the start time 1998-01-30T11:30 and t₂ is ∞ (an open-ended event). More elaborate taggings are allowed in the TIMEX2 formalism:

(6) (i) <TIMEX2 VAL="PRESENT_REF" ANCHOR_VAL="1998-01-30T11:30" ANCHOR_DIR="AS_OF">
 Right now
 </TIMEX2>
 (ii) <TIMEX2 VAL="2000-10-27">
 tomorrow
 </TIMEX2>
 (iii) <TIMEX2 VAL="PT30M">
 half an hour long
 </TIMEX2>

(iv) <TIMEX2 VAL="1990" MOD="BEFORE">
 more than a decade ago
 </TIMEX2>

For the remainder of this paper, we use an abbreviated version of the notation above, e.g., 11:30 instead of 1998-01-30T11:30. We focus on the mapping of the time-stamped input into a *matrix* (i.e., main) clause and an *adjunct* (i.e., subordinate) clause conjoined by a connecting word. Consider the following input form:

(7) fall(John,15:01,15:01) \wedge laugh(Mary,15:01,15:03)

This logical expression may be expressed in several different matrix/adjunct combinations including *Mary laughed while John fell*, *Mary laughed after John had fallen*, *Mary had laughed as John fell*, but not, for example, *Mary laughed before John fell*, *Mary was laughing until John fell*, etc.³ When the facts are expressed in the same sentence, aspectual considerations and the choice of connecting words become important.

A major contribution of our work is the development of a relationship between Allen's temporal intervals (1983; 1984) and Hornstein's (1990) theory of tense—and the implementation of an algorithm that uses this relationship for generating a matrix/adjunct sentence. This relationship was previously investigated by Yip (1985) who extended both frameworks to handle particular aspectual features. However, Yip's work omitted the perfect aspect, treated all event verbs as dynamic (i.e., non-stative), and attempted to handle all dynamic verbs with a single set of temporal constraints.

Brent (1990) also investigated an integration of tense theories by Hornstein and Allen; however, this approach considered only punctual events—events with no duration—and lexical connectives were treated syntactically rather than semantically. This approach produced spurious temporal assignments such as $E_1=E_2$ for the temporal interpretation of sentences such as *John went to the store before Mary arrived* (where E_1 is the time of *John went to the store* and E_2 is the time of *Mary arrived*). Neither Brent nor Yip took telicity (i.e., event culmination) or atomicity (i.e., event duration) into account in their analyses. Furthermore, neither approach has been applied to generation.

In our framework, events are allowed to have duration and are viewed in terms of a fuller theory of aspect through the use of Allen's interval theory. We show how constraints on aspect affect the final selection of aspectual features; and we analyze how aspectual selections can alter the meanings of connecting words and thus affect their final selection. We illustrate the algorithm by showing the full set of sentences to which our linguistic constraints are applied.

Our approach uses a standard AI technique of constraint compilation and table look-up, which eliminates most—but not all—of the overgeneration. While producing multiple possibilities may seem problematic, we note that it is not

³ For ease of exposition, we take the first literal to be the matrix clause and the second literal to be the adjunct clause. However, in our approach, these may be re-ordered, e.g., if a sentence cannot be generated for the first ordering.

our goal to provide *the* one choice of natural language sentence for event combinations, but to provide a set of ranked sentences that are legal (both temporally and grammatically). Ultimately, these possibilities may be passed as input to a full text planning system, e.g., (Hovy, 1993). However, we view CONGEN as a self-contained module designed to enhance the performance of a full MT system, where overgeneration is constrained by linguistically motivated rules and statistical extraction techniques applied to a word lattice of possibilities (Habash and Dorr, 2002).

Other inferencing capabilities may also be applied to our work. The framework of Amsili and Rossari (1998) uses sophisticated causal information to produce an appropriate connective. In other work (Dale, 1992; Reiter and Dale, 1997), referring expressions are induced from natural-language databases. The paradigm of Danlos (1999; 2000) uses extra linguistic-knowledge (causality) to determine the ordering of sentential clauses. Techniques developed by Elhadad and McKeown (1990) for a “deep generator” are designed to select temporal connectives from pragmatic features. The framework of Gagnon and Lapalme (1996) generates temporal information from an incrementally updated conceptual representation (Discourse Representation Structure).

Such approaches may provide important knowledge supporting the generation of temporally related clauses, but they also impose added complexity. In the absence of these more knowledge-intensive (often multi-sentence) approaches, we make use of a “simplicity heuristic” that allows us to produce a single sentence from an underlying form. It is expected that the replacement of this heuristic with more pragmatic text planning techniques would be one possible area for future investigation.

3 Background

Both aspectual and temporal knowledge are used for generation of natural language expressions that reflect temporal relations present in underlying concepts. This section distinguishes the representations used for these two types of knowledge.

3.1 Aspectual Knowledge

Following Dowty (1979) and Vendler (1967), aspect is taken to have two components, one comprises *non-inherent* features (e.g., those features that define the perspective such as simple, progressive, and perfect) and another comprises *inherent* features (e.g., those features that distinguish between states and events). Non-inherent features are dependent on temporal context; thus, they are not stored with the lexical item and may be controlled during language generation. These are distinguished from inherent features, which are stored with the lexical item and are used for lexical selection.

Suppose we are generating a sentence from the following time-stamped input:

(8) go(John,store,15:00,15:15) \wedge arrive(Mary,15:31,15:31)

For the purpose of illustration, assume that the current time is 18:00. The two literals in (8) may be realized in a number of different aspectual combinations:⁴

- (9) (i) John went to the store before Mary arrived. (simple/simple)
- (ii) John went to the store before Mary had arrived. (simple/perfect)
- (iii) John had gone to the store before Mary arrived. (perfect/simple)
- (iv) John had gone to the store before Mary had arrived. (perfect/perfect)

The aspectual variations shown here are primarily a function of values of non-inherent features (i.e., perfect *vs.* simple). These feature values must be determined before the two events can be combined since this information is necessary for selecting the appropriate temporal connectives (e.g., *after*, *before*, *when*, *while*, etc.).

A number of aspectually oriented representations have been proposed for inherent features that readily accommodate the types of aspectual distinctions that are of concern here (Bach and Harms, 1968; Comrie, 1976; Crouch and Pulman, 1993; Dowty, 1979; Hwang and Schubert, 1994; Jackendoff, 1983; Jackendoff, 1990; Mourelatos, 1981; Nirenburg and Pustejovsky, 1988; Olsen, 1994; Passoneau, 1988; Pustejovsky, 1988; Pustejovsky, 1991a; Pustejovsky, 1991b; Pustejovsky, Bergler, and Anick, 1993; Steedman, 1997; Vendler, 1967). The current model implements an aspectual classification using three features proposed by Bennett et al. (1990) following the framework of Moens and Steedman (1988): [\pm dynamic] (i.e., events *vs.* states), [\pm telic] (i.e., culminative events (transitions) *vs.* nonculminative events (activities)), and [\pm atomic] (i.e., point events *vs.* extended events).⁵

Consider the two verbs *ransack* and *obliterate*. These are distinguished by means of aspectual features: [+d,-t,-a] for the verb *ransack* and [+d,+t,+a] for the verb *obliterate*. Although these two verbs are semantically similar, the feature-based framework accounts for surface distinctions such as the following:

- (10) (i) John ransacked the house every day.
- (ii)* John obliterated the house every day.

To summarize, we adopt the feature-based scheme of Bennett et al. (1990) and Moens and Steedman (1988), but we have found this to be compatible with the schemes used by several researchers, as illustrated in Table 1. We will refer primarily to the feature notation in column 1 of this table, but we will use this interchangeably with aspectual labels in column 2 (state, activity, achievement, accomplishment) and column 3 (state, process, event), depending on the level of specificity required by the discussion.

⁴ The term *perfect* refers to either the present or the past (plu) perfect (i.e., it does not specify the tense).

⁵ Androutsopoulos (1999) provides a coarser-grained analysis of aspectual categories of verbs, classified in a specific (airport) domain. Our goal is to provide a wide range of verbs, taking into account values of durativity, telicity, and atomicity—for broader applicability.

Table 1. Classification of Inherent Aspectual Features

B+90,M88	D79,V67,P88	M81,C76,B86	N&P87,P88	OI94	J83,J90	Example
[-d]	state	state	state	state**	state(BE)	be, like, know
[+d,-t,-a]	act(ext)	process	process	act	event(GO)	ransack, run
[+d,-t,+a]	act(pt)	process	trans*	semelfactive	event(GO)	tap, wink
[+d,+t,+a]	ach	event	culm trans	ach	event(GO)	obliterate, win
[+d,+t,-a]	acc	event	log trans	acc	event(GO)	destroy, arrive

B+90=[Bennett et al., 1990], M88=[Moens&Steedman, 1988], D79=[Dowty, 1979], V67=[Vendler, 1967], P88=[Passonneau, 1988], M81=[Mourelatos, 1981], C76=[Comrie, 1976], B86=[Bach, 1986], N&P87=[Nirenburg&Pustejovsky, 1987], OI94=[Olsen, 1994], J83=[Jackendoff, 198], J90=[Jackendoff, 1990]

*Transitions are further classified as simple (give) and causative (send)

**States are further classified as individual (know) and stage level (be sick)

3.2 Temporal Knowledge

Tense is taken to be the external time relationship between a given situation and others. (See, for example, (Bennett et al., 1990)). For past, present and future tenses, the relationship between the time of an event being discussed, say E, and speech time, say S, may be characterized as $E=S$ for present tense; $E<S$ for past tense; and $S<E$ for future tense. However, as one considers tenses like past perfect (e.g., *Joe had left the office by 3:00pm*), present perfect (e.g., *Mary has left the office*), and future perfect (e.g., *John will have left the office by 3:00pm*), speech time and event time are insufficient to distinguish the properties of the tenses. For example, in both the past perfect tense and the past tense, event time precedes speech time, that is, $E<S$.

Reichenbach (1947) noticed that the introduction of a reference timepoint, which he labeled R, provides enough information to characterize all of the tenses that occur in natural language (Hornstein, 1990). To illustrate R, consider the past perfect tense. There is some point in time that occurs between event time and speech time. Prior to this intermediate point in time, the event being described has already occurred; thus, for the past perfect tense, we may say that $E<R<S$.

More formally, the Reichenbachian framework for tense postulates three theoretical entities: S (the moment of speech), R (a reference point), and E (the moment of the event). The key idea is that certain linear orderings of the three time points get grammaticalized into basic tenses. English uses six basic tenses with the following Basic Tense Structures (BTS):

S,R,E	present, present progressive
E,R,S	past, past progressive
S,R,E	future, future progressive
E,S,R	present perfect
E,R,S	past perfect
S,E,R	future perfect

The S, R, and E points may be separated by a line (in which case, the leftmost point is interpreted as temporally earlier than the other) or by a comma (in which case, the points are interpreted as contemporaneous).

The approach adopted here is based on a neo-Reichenbachian framework proposed by Hornstein (1990) in which two BTSs are organized into a Complex Tense Structure (CTS) as follows: the BTS of the first (matrix) clause is written over the BTS of the second (adjunct) clause and the S and R points are then associated. In the general case, the association of the S and R points may force the R point in the second BTS (R_2) to be moved so that it is aligned with the R point in the first BTS (R_1). (In the current example, the two R points are already aligned.) The second E point (E_2) is then placed accordingly. For example, the CTS for the sentence *John went to the store before Mary arrived* would be specified as follows:

$$\begin{array}{ccc} E_1, R_1 & \text{---} & S_1 \\ | & & | \\ E_2, R_2 & \text{---} & S_2 \end{array}$$


We now return to example (8). If we take the current time (S) to be 18:00, then $E < S$ for both literals, i.e., there are three possible BTSs corresponding to each literal: E, R, S (e.g., went, was going), E, R, S (e.g., has gone), and E, R, S (e.g., had gone). The multiplicative combination of pairwise BTSs (three per literal) yields 9 possible CTSs, some of which correspond to ungrammatical sentences:⁶

- (11) (i) John went (was going) to the store before Mary arrived (was arriving).
- (ii)* John went (was going) to the store before Mary has arrived.
- (iii) John went (was going) to the store before Mary had arrived.
- (iv)* John has gone to the store before Mary arrived (was arriving).
- (v) John has gone to the store before Mary has arrived.
- (vi)* John has gone to the store before Mary had arrived.
- (vii) John had gone to the store before Mary arrived (was arriving).
- (viii)* John had gone to the store before Mary has arrived.
- (ix) John had gone to the store before Mary had arrived.


The focus of our generation task is on the elimination of illegal clause combinations and the selection of a connecting word (e.g., *before*).

Tense is determined by factors relating not to the particular lexical tokens of the surface sentence, but to the temporal features of the context surrounding the event coupled with certain linguistically motivated constraints on the tense structure of the sentence. In particular, it has been argued persuasively by Hornstein (1990) that all sentences containing a matrix and adjunct clause are subject to a linguistic (syntactic) constraint on tense structure *regardless* of the lexical tokens included in the sentence. For example, Hornstein's linguistic Constraint on Derived Tense Structures (CDTS) requires that the association of S and R points not involve crossover in a complex tense structure:

⁶ The nine CTSs correspond to 16 possible sentences if one considers the progressive/non-progressive distinction.



This structure would be associated with sentence (11)(ii) above: * *John went (was going) to the store before Mary has arrived*. Here, the association of R_2 and R_1 violates the CDTs, thus ruling out the sentence. Sentences in (11)(iv), (vi), and (viii) are also ruled out by the CDTs, as shown here:



Note that this linguistic constraint is a syntactic restriction on the manipulation of tense structures, not on the temporal interpretation of tensed sentences. Thus, the constraint holds regardless of the lexical token that is chosen as the connecting word between the two events:

(12)

$$* \text{ John went to the store } \left\{ \begin{array}{l} \text{as} \\ \text{before} \\ \text{after} \\ \text{as soon as} \\ \text{while} \end{array} \right\} \text{ Mary arrives.}$$

Hornstein's theory crucially relies on an asymmetry between the matrix and adjunct clauses. Thus, there is an important distinction between the work of Hornstein (1990), in which the asymmetrical property is fundamental to the theory, and that of Yip (1985) in which the asymmetrical property is entirely abandoned. Hornstein's intuition is the one adopted here based on the observation that we cannot arbitrarily interchange the matrix and adjunct clauses. For example, Yip's theory predicts that we should be able to replace "X after Y" with "Y before X," which is not always the case:

(13) (i) John will go to the store after Mary has arrived.

(ii)* Mary has arrived before John will go to the store.

(14) (i) John will go to the store after Mary arrives.

(ii)* Mary arrives before John will go to the store.


Given this asymmetrical property, it would not be possible to randomly select a matrix/adjunct order and an appropriate temporal connective for a surface sentence solely on the basis of lexical information. What is needed is the temporal relation between the two events and the constraints on their combination before it is possible to derive the matrix/adjunct ordering of the sentence.

4 Handling Events with Duration

Hornstein’s (1990) theory of tense assumes that events are points in time. To extend this theory to events that have duration, we analyze events in terms of Allen’s (1983; 1984) theory of temporal interval relationships, which has been used for a number of artificial intelligence and natural language understanding applications (Allen, 1983; Galton, 1990; Lesperance and Levesque, 1990; Vilain, Kautz, and van Beek, 1990; Williams, 1990).

Allen proposes the existence of seven basic relationships and their inverses between two intervals: *before* ($<$), *after* ($>$) *during* (d), *contains* (di), *overlaps* (o), *overlapped by* (oi), *meets* (m), *met by* (mi), *starts* (s), *started by* (si), *finishes* (f), *finished by* (fi), and *equal* ($=$).⁷ These relationships are illustrated in Table 2.

Table 2. Allen’s 13 Interval Relationships



In this figure, a line indicates an interval. The intervals shown here are closed, as indicated by closed circles at either end. Shortly, we will see examples of open intervals, which would be indicated by open circles.

To associate a tense with an event that has duration, we first determine the interval relationship between the event time interval and speech time. A BTS is initially associated with the event if it preserves the relationship between the event time E and speech time S . For example, if it is determined from a logical expression that the event *John went to the store* occurs before the speech time, three possible BTSs for the event are: E, R_S (past), E_S, R (present perfect),

⁷ The inverse of *equal* is *equal*, so there are a total of 13 different interval relationships.

and E_R_S (past perfect). In each case, at least one line separates event time E and speech time S , indicating that E occurs before S .

Note that there are three BTS possibilities because the R (reference) point is not yet established. R is not part of the static knowledge associated with an event; rather, R emerges from the operation of relating one event with another. A given BTS may later be eliminated as a possibility for an event if the relation between that BTS and another reveals an inconsistency between the two R points (i.e., a violation of the CDTS). Thus, the R point does not serve to narrow down the BTS possibilities for a given literal, but it serves to reduce the space of possibilities for *combining* two literals into a matrix and adjunct clause, by virtue of its role in the CDTS.


The full extension of Hornstein’s theory to events with duration requires a more detailed analysis of the E point in the BTS representation. In particular, we require E to be divided into a start time E_s and a stop time E_f , corresponding to the time-stamps in the logical expression. We shall denote the interval as E_{sf} . A second interval (actually a point) is defined as the current (speech) time denoted by S . The time interval for a literal may be *open* (corresponding to a stop time of ∞) or *closed* (corresponding to a stop time containing an actual value). Given a time-stamped logical expression and the current time, we can obtain a partial ordering over E_s , E_f , and S , and we can derive the temporal interval relationship between E_{sf} and S with Allen’s representation.

Table 3 represents the full extension of Hornstein’s BTS representation to events that have duration. Appendix A contains our analysis that led to this table of each possible alignment of E_s , E_f , and S for open and closed event time-stamps. The table shows the mapping from events that are either points or intervals into BTSs. From the partial ordering and the interval relationship, we have determined all allowable BTSs for closed intervals, points, and open intervals for each possible relationship of the interval to S . The last three cases in Table 3 cover Hornstein’s original analysis.

Consider the previous example (8), repeated here as (15):

(15) $go(john,store,15:00,15:15) \wedge arrive(mary,15:31,15:31)$

Let the label ES_1 refer to the E/S relationship for the first literal, and let the label ES_2 refer to the E/S relationship for the second literal. If the speech time (S) is 18:00, then ES_1 is represented as a closed interval preceding S and ES_2 is represented as a point interval preceding S :



The first E/S relationship corresponds to the first case in Table 3 since the entire closed interval event precedes the speech time. The second E/S relationship corresponds to the ninth case in Table 3 since the point interval precedes the speech time. Each E/S relationship is associated with three BTSs: past tense (E_R_S); past perfect (E_R_S); and present perfect (E_S_R). All of these preserve the ordering between E_s and S and between E_f and S . Hornstein’s CDTS

Table 3. Mapping Between E/S Time Relationships and Allowable BTS's

Time Points	Salient Relationship	Allowable BTSs
E_s E_f S 	$E_s < E_f < S$	E_{sf}, R_S (past) $E_{sf}_R_S$ (past perf.) E_{sf}_R, S (pres. perf.)
S E_s E_f 	$S < E_s < E_f$	S_R, E_{sf} (fut.) $S_E_{sf}_R$ (fut. perf.) S, R, E_{sf} (pres.)
E_s S E_f 	$E_s < S$	E_s, R_S (past) $E_s_R_S$ (past perf.) E_s_R, S (pres. perf.)
	$S < E_f$	S_R, E_f (fut.) $S_E_f_R$ (fut. perf.)
E_s E_f S 	$E_s = S$	S, R, E_s (pres.)
	$S < E_f$	S_R, E_f (fut.) $S_E_f_R$ (fut. perf.)
	$S = E_s < E_f$	S, R, E_{sf} (pres.) S_R, E_{sf} (fut.) $S_E_{sf}_R$ (fut. perf.)
E_s E_f S 	$E_s < S$	E_s, R_S (past) $E_s_R_S$ (past perf.) E_s_R, S (pres. perf.)
	$E_f = S$	S, R, E_f (pres.)
	$E_s < E_f = S$	S, R, E_{sf} (pres.) E_{sf}, R_S (past) E_{sf}_R, S (pres. perf.)
E_s S E_f 	$E_s < S$	E_s, R_S (past) $E_s_R_S$ (past perf.) E_s_R, S (pres. perf.)
E_s E_f S 	$E_s = S < E_f$	E_s, R, S (pres.)
S E_s E_f 	$S < E_s$	S_R, E_s (fut.) $S_E_s_R$ (fut. perf.)
	$S < E_f$	S_R, E_f (fut.) $S_E_f_R$ (fut. perf.)
E_{sf} S 	$E_{sf} < S$	E_{sf}, R_S (past) $E_{sf}_R_S$ (past perf.) E_{sf}_R, S (pres. perf.)
E_{sf} S 	$E_{sf} = S$	S, R, E_{sf} (pres.)
S E_{sf} 	$S < E_{sf}$	S_R, E_{sf} (fut.) $S_E_{sf}_R$ (fut. perf.)

(described above in Section 3) can be used to identify which pairs of BTSs for the two literals are allowed to occur together in a complex matrix/adjunct sentence.

In the next section we will describe an algorithm that includes the application of the CDTS for the realization of tense, aspect, and connecting words for two literals. We will show that this algorithm relies on the temporal relationship between the two literals to assign possible BTSs and, ultimately, to produce legal BTS combinations.

5 Algorithm for Selection of Tense, Aspect, and Connecting Words

In this section, we present the algorithm implemented in CONGEN for generating matrix/adjunct structures from conjunctions of literals. The input to the algorithm is a conjunction of two time-stamped literals and their corresponding verb tokens; these verbs are associated with inherent aspectual features provided by the lexicon of the associated application. The algorithm seeks to place the verb tokens in a matrix/adjunct structure if possible. If there are several allowable realizations for a given conjunction, then all alternatives are produced—in a ranked ordering. Thus, the output of the algorithm is a ranked list of legal temporal connectives C, BST pairs P, and aspect A_1 and A_2 . The end application may seek to vary its ultimate choice of surface realization based on this ranking coupled with the constraints of the application domain.

Figure 2 shows the overall algorithm for generating matrix-adjunct pairs. As mentioned previously (in fn. 3), we allow for a reordering of the matrix and adjunct clauses, e.g., if a sentence cannot be generated for the first ordering. We implement this by applying this algorithm twice in such cases.

There are two main steps in this algorithm: Selection of tense and aspect, with sub-steps a–d; and Selection of a connecting word, with sub-steps e–g. Step a is a straightforward application of the framework described in Section 4 (Table 3). The remaining steps require elaboration; we will briefly describe each of these steps in turn below.

Note that the ordering of steps in Figure 2 does not matter as much for completeness as for efficiency. It is generally advantageous to apply linguistic constraints as soon as possible. When tense is selected before aspect, the CDTS may be applied immediately to eliminate illicit tenses; the alternative order would require the CDTS to be applied after aspect selection has already multiplied out many illicit possibilities. In addition, there are some data dependencies in the algorithm, e.g., step d (selecting between progressive and simple aspect) requires that the tense (BTS) already be established.

5.1 Tense Selection Process

As we saw in above, BTSs are determined for each event in the logical expression based on the interval relationship between event time and speech time. This is step a of the algorithm in Figure 2. Step b, the tense selection process, must then

Input:

- A **query** consisting of time-stamped literals L_1 and L_2 and verbs V_1 and V_2 for events E_1 and E_2
- A **lexical specification** for verbs V_1 and V_2 consisting of inherent verb features F_1 and F_2 , where $F_i = (\pm\text{atomic}, \pm\text{dynamic}, \pm\text{telic})$.

Output:

- A **ranked list of BTS pairs** P .
- A **list of aspectual perspective pairs** A for each BTS pair in P .
- A **list of ranked sets of temporal connectives** C for each BTS pair in P .

Procedure:

- I. Select tense and aspect
 - a. Find allowable BTS sets, BTS_1 and BTS_2 , for L_1 and L_2 based on time point relationships. (Section 4—Table 3)
 - b. Apply CDTS on $BTS_1 \times BTS_2$ to obtain a list of allowable matrix/adjunct BTS pairs, P . (Section 5.1—Table 4) NOTE: *Perfect* vs *non-perfect* emerges from the BST.
 - c. Sort P by frequency based on corpus analysis. (Section 7—Table 10)
 - d. Using the inherent lexical features F_1 and F_2 , find one aspectual perspective pair (*progressive* vs. *simple*) for each BTS pair in P —producing A . (Section 5.2—Figure 4)
- II. Select connecting word
 - e. Determine interval relation $T \in \{<, >, =, d, di, o, oi, m, mi, s, si\}$ between L_1 and L_2 . (Section 4—Table 2)
 - f. For each BTS pair in P , find the list of allowable temporal connective sets C , ranking choices within each set according to sparseness (“simplicity heuristic”) based on the following:
 - i. Interval relation T
 - ii. Aspectual features:
 - Inherent features F_1 and F_2 ($\pm\text{atomic}, \pm\text{dynamic}, \pm\text{telic}$)
 - Perspective pair from A (*progressive* vs. *simple*) that corresponds to the BTS pair
 - g. For each BTS pair in P , sort the associated connective set in C by the frequency based on corpus analysis. (Section 7—Fully expanded version of Table 13)

Fig. 2. Algorithm Behind CONGEN: Producing Matrix/Adjunct Sentences Reflecting Temporal Relations

determine which combinations of BTS pairs are legal using the CDTS outlined in Section 3. Any tense pairs that do not contain a crossover in the corresponding complex tense structure may be used as a possible tense in a complex sentence.

We have precompiled the allowable tense pairs by combining each basic tense with every other basic tense and then ruling out those that are disallowed by the CDTS. This precompilation procedure produced the table of allowable tense pairs shown in Table 4. Here, each tense in the left-hand column may be legally paired with each tense in the right-hand column. The tense selection process ensures that any BTS pair that is not in this table is eliminated. Thus, as alluded to in Section 4, the R (reference) point in each BTS is indirectly accessible, by virtue of its role in the CDTS which provides the basis for this table.

Table 4. Allowable Tense Pairs for Matrix/Adjunct Sentences (CDTS)

Future Tenses:

<i>Matrix Tense</i>	<i>Adjunct Tense</i>
Future	Present
Future Perfect	Present Perfect
	Future
	Future Perfect

Past Tenses:

<i>Matrix Tense</i>	<i>Adjunct Tense</i>
Past	Past
Past Perfect	Past Perfect

Present Tenses:

<i>Matrix Tense</i>	<i>Adjunct Tense</i>
Present	Present
Present Perfect	Present Perfect

Reconsider the conjunction given earlier in (15): $\text{go}(\text{john}, \text{store}, 15:00, 15:15) \wedge \text{arrive}(\text{mary}, 15:31, 15:31)$. Recall that the time of speech was 18:00. From the ES1 and ES2 relationships in (16), we determine that the set of allowable BTS pairs for each literal is {past, past perfect, present perfect}.⁸ Suppose that the first literal has been selected as the matrix. Then for each of the three basic tenses for the matrix literal, we use the table of allowable tense pairs, compiled from the CDTS, to determine the allowable adjunct tenses. The resulting matrix/adjunct pairs (the value of P in step b of Figure 2) are the following: {(past,past),(past,past perfect),(past perfect,past), (past perfect,past perfect), (present perfect, present perfect)}.

⁸ For clarity, the BTS representations are shown here in the more readable form, e.g., (past,past) instead of (E,R_S,E,R_S).

Once tenses are selected for time-stamped literals, there are frequently many legal combinations. In step c of Figure 2, a corpus-based ranking is applied to the possible tense combinations.⁹ For example, when both *perfect* and *non-perfect* are compatible with the CDTS, both alternatives are produced but the simpler (non-perfect) form is generally ranked higher. While this may seem like an unprincipled choice, our own large-scale corpus analysis reveals that, indeed, the *simple* aspect is 8–10 times more likely to be selected than the *perfect* aspect. Section 7 describes this corpus analysis in more detail. The result of the ranking is the following list: {(past,past),(past,past perfect),(past perfect,past), (present perfect, present perfect), (past perfect,past perfect)}.

Note that the last two BTS pairs have swapped places with respect to the unranked list given above.

For the purpose of illustration, suppose that the temporal connecting word *before* is to be selected (by an independent process) to connect the two sentences.¹⁰ We can then generate the following alternative sentences (given sufficient grammatical information about the two literals):

- (17) (i) John went / was going to the store before Mary arrived / was arriving.
- (ii) John went / was going to the store before Mary had arrived.
- (iii) John had gone to the store before Mary arrived / was arriving.
- (iv) John has gone to the store before Mary has arrived.
- (v) John had gone to the store before Mary had arrived.

Next, we shall see how the choice of non-inherent aspectual features (e.g., *simple vs. progressive*) can be narrowed down using the temporal interval information. Then, in Section 5.3, we show how the selection of the connecting word interacts with the final selection of the tense and aspectual features.

5.2 Aspect Selection Process

As described in Section 3.1, aspect is taken to have two components, one comprised of *non-inherent* features and another comprised of *inherent* features. Since inherent features are lexically specified, they are fixed; thus, in generation, selecting aspect involves finding values for non-inherent features (step d of the algorithm in Figure 2). The two aspectual features that are not inherent are: (1) *perfect vs. non-perfect* and (2) *progressive vs. simple*. Together these two features define the *perspective* of a verb phrase. The final aspectual realization of the generated sentence emerges from the composition of inherent verb properties and these chosen values.

As noted above, when both *perfect* and *non-perfect* are compatible with the CDTS, both alternatives are produced but the simpler (non-perfect) form is

⁹ In a future version of the system, we will also incorporate the tense-selection scheme of (Olsen et al., 2000; Olsen et al., 2001), where aspectual attributes such as telicity are used to further constrain the choices of possible tenses.

¹⁰ As we will see shortly, the choice of tense and aspectual features is further narrowed down before the temporal connective is selected. The inclusion of the temporal connective *before* in (17) is for illustrative purposes only.

generally ranked higher (at the end of step c in Figure 2) if it is still available as a legal possibility. However, there are often cases where the perfect aspect will be the *only* possible realization for a given pair of time-stamped literals, depending on which temporal connective is selected in a later step. For example, the connective *when* would require the perfect aspect in order to adequately capture the event ordering in the sentence *John will have gone when Mary arrives*, whereas the connective *before* could convey this same information with or without the perfect aspect: *John will have gone before Mary arrives* or *John will go before Mary arrives*. If *when* were chosen, the future perfect would be generated, whereas if *before* were chosen, the simple future tense would be generated. Thus, the choice of perfect vs. non-perfect will be further constrained in step f of the process, when the connective is selected. However, if both *perfect* and *non-perfect* are available simultaneously, the *non-perfect* is generally preferred.

Our method for selecting between *progressive* and *simple* (step d in Figure 2) relies on a set of restrictions based on the work of Dowty (1979) and Olsen (1997) adapted for generation of temporal information. We have recast Dowty’s constraints on the relationship between inherent verb features and the progressive/simple choice in terms of a decision tree, as illustrated in Figure 3. This decision tree maps directly into the algorithm shown in Figure 4. All inherent features (\pm atomic, \pm dynamic, \pm telic) are taken from an English lexicon of 10,000 verb entries that has been made publicly available for research purposes (Dorr, 2001).

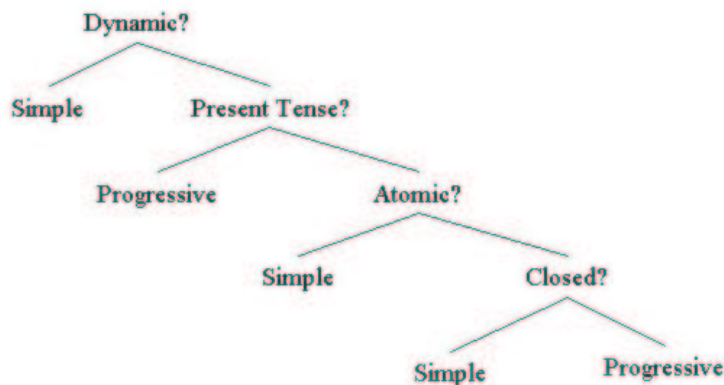


Fig. 3. Recasting Dowty’s Constraints as a Decision Tree

Constraint I is consistent with Dowty’s constraint on states [-dynamic], i.e., that a stative verb cannot participate in the progressive construction: **I was knowing the answer*.

Input:

- A **time-stamped literal** L
- A **verb** V (corresponding to L)
- A **BTS** B (provided by step a of Figure 2).

Output: An aspectual perspective for V.

Procedure: Choose aspectual perspective of V subject to the following constraints:

- I. If V is inherently a state [-dynamic], then V’s aspectual perspective is *simple*.
- II. If B corresponds to the present tense (S,R,E), then V’s aspectual perspective is *progressive*.
- III. If the interval for L is a point, that is, the start time and stop time are the same, and V is [+atomic], then V’s aspectual perspective is *simple*.
- IV. If the interval for L is not a point, then: (a) If the interval is closed (complete), then V’s perspective is *simple* aspect; (b) if the interval is open (incomplete), then V’s aspectual perspective is *progressive*.

Fig.4. Adaptation of Dowty’s Constraints: Algorithm for Selecting between Simple and Progressive

Constraint II is the only one that relies on knowledge of the tense selected for the natural language verb (i.e., the result of steps a–c of the algorithm in Figure 2). A present tense verb has only one realization possibility, progressive, because the simple present tense is typically used to characterize some other component of meaning such as habitual or generic action: *He brushes his teeth (regularly)*.¹¹

Constraint III refers to the case of a [+atomic] verb (such as *cough*), where the instantaneous time interval eliminates the possibility of using the progressive. Note that a point-interval literal can never be associated with a [-atomic] verb; if a [-atomic] verb (e.g., *run*) is selected for a point interval, no natural-language realization will be produced. Recent work on aspect (Olsen, 1997) provides a *privative* analysis that predicts this non-correspondence. In this approach, the marked feature [+durative] cannot be changed to its unmarked counterpart [\emptyset durative].¹²

Constraint IV (a) allows for the realization of a complete action as a simple verb (e.g., *John won/ran*) while IV (b) allows an incomplete action to be realized as a progressive verb (e.g., *John was winning/running*). This constraint provides the mechanism necessary for covering the cases of coercion discussed by Dorr (1992), where an inherently [+atomic] verb (e.g., *win*) is interpreted as [-atomic] in a particular context. Olsen’s privative analysis provides a more systematic account, where an achievement is inherently unmarked for durativity ([\emptyset durative]) but can become marked ([+durative]), behaving more like an activ-

¹¹ This point is noted in (Hornstein, 1990, p. 206, fn.20), where the present tense is shown many interpretations that are nontemporal.

¹² Clearly the selection of the natural-language verb is dependent on both temporal and aspectual information in such cases; this point is discussed further in Section 8.

ity in the context of a progressive morpheme (“-ing” in English). Constraint IV provides the necessary mechanism for retaining the non-point interval reading, i.e., the desired result for either analysis, coercion or privative.

Note that constraint IV subsumes the handling of telic verbs which, when realized in the past progressive (i.e., *was winning*), convey the notion that the action is not necessarily complete (as in *has won*). Thus, the handling of all three types of inherent features are covered by the algorithm: constraint I accounts for $[\pm\text{dynamic}]$; constraint III accounts for $[\pm\text{atomic}]$; and constraint IV accounts for $[\pm\text{telic}]$.

In our ongoing example, recall that the CDTS has pared down the sentences down to the 5 remaining cases in (17)—or 10 possibilities if the progressive is included. Note that both literals are associated with closed, past temporal intervals. The verb *go* corresponds to a non-point interval. The verb *arrive* corresponds to a point interval and is inherently $[-\text{atomic}]$. Following the algorithm in Figure 4, constraint IV (a) induces the selection of simple for *go* and constraint III induces the selection of simple for *arrive*. Thus, neither the matrix nor the adjunct is generated in the progressive form and the resulting value of A (in step d of Figure 2) is: $\{(\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), \}$. Thus, the sentences are pared down to the following five cases:

- (18) (i) John went to the store *CW* Mary arrived.
- (ii) John went to the store *CW* Mary had arrived.
- (iii) John had gone to the store *CW* Mary arrived.
- (iv) John has gone to the store *CW* Mary has arrived.
- (v) John had gone to the store *CW* Mary had arrived.

Note that we are using *CW* as a “connecting word” placeholder since, at this point, the temporal connective has not yet been selected.

Ultimately, after the temporal connective is chosen, the first of these five cases will still be ranked highest. As we will see in Section 5.3, the final ranking is not brought about through application of Hornstein’s CDTS, but by a “simplicity heuristic” and a corpus-based ranking that is applied during the selection of the temporal connective. This is an important contribution of our work: We assign a low ranking to certain semantically anomalous cases (e.g., *John has gone to the store before Mary has arrived*) even when such cases are syntactically permissible in Hornstein’s framework.

To further illustrate each of the constraints in aspect selection, consider the following literals:

- (19) (i) $\text{red}(\text{apple1}, t_1, t_2)$
- (ii) $\text{cough}(\text{john}, t_3, t_4)$
- (iii) $\text{win}(\text{john}, \text{race}, t_5, t_6)$
- (iv) $\text{walk}(\text{john}, \text{home}, t_7, t_8)$

Suppose that, in all four cases above, the temporal interval precedes the speech time (S). In addition, suppose that only case (19) (ii) involves a point interval and only case (19) (iii) involves an open interval.

For (19)(i), the literal corresponds to the stative [-dynamic] verb *be (red)*. The sentences *the apple was being red* and *the apple was red* both convey the same information, but constraint I of Figure 4 eliminates the first sentence as a possibility and the simple form is produced.

Alternatively, consider (19)(ii) which describes a point activity [+atomic], *cough*. The sentences *John coughed* and *John was coughing* both convey that the activity ended, but only the first sentence corresponds to a point interval. Thus, constraint III induces the selection of the simple form.

The literal in (19)(iii) describes an incomplete action *win*. The sentences *John won the race* and *John was winning the race* both convey the meaning of this action, but only the second case indicates that the final state of completion may not have been achieved. Thus, constraint IV (a) induces the selection of the progressive form.

In contrast, the literal in (19)(iv) describes a completed action, *walk*. The sentences *John walked home* and *John was walking home* both convey that this activity is extended. However, only the first case indicates that the final state of completion has been achieved. Thus, constraint IV (b) induces the selection of the simple form.

5.3 Selecting Temporal Connecting Words

In our ongoing example, we assumed the temporal connective between the two sentential concepts would be selected by an independent process. In this section, we discuss this process, i.e., steps e–g of Figure 2. The list of temporal connectives for this research was extracted from the most frequently occurring cases in the entire Lancaster-Oslo-Bergen (LOB) tagged corpus of 53,411 sentences, namely, *after*, *before*, *since*, *until*, *when*, and *while*. (Our corpus analysis is described in more detail in Section 7.)

Two pieces of information contribute to the selection of a temporal connecting word for a matrix/adjunct pair in step f of Figure 2: the temporal interval relationship between the two literals provides a means to select a particular subset of candidate connecting words. This is established in step e of the algorithm (using Table 2 of Section 3.2). In our current example, the temporal interval T is determined from (16) to be *before* (<). Second, in step f, the aspectual features are used to further restrict the set of possible connecting words for each BTS pair; this involves inherent features F₁ and F₂ (e.g., [+dynamic] *vs.* [-dynamic]) and non-inherent perspective pairs in A (i.e., progressive *vs.* simple).

With respect to inherent features, we discovered in our analysis that the aspectual distinction most relevant to the choice of temporal connective is the state/event distinction (i.e., [-dynamic] *vs.* [+dynamic])—a distinction that is readily extractable from our English lexicon (Dorr, 2001). By contrast, the distinction between activities and achievements or between activities and accomplishments has no impact on the choice of temporal connective. Thus, we broadly classify verbs in terms of S=[-dynamic] and D=[+dynamic], following the state/event distinction suggested by (Jackendoff, 1983; Jackendoff, 1990)

(column 6 of Table 1). With respect to non-inherent features, we focus on the progressive/simple distinction (p=progressive and s=simple). We shall abbreviate [+dynamic]/[+progressive] as Dp; [+dynamic]/[-progressive] as Ds; and [-dynamic]/[-progressive] as Ss (since [-dynamic] is a state).¹³

A more fine-grained analysis shows that the perfect/non-perfect distinction also has an impact on the choice of temporal connective (e.g., the choice of *when* over other connectives as mentioned earlier in Section 5.2). The selection of a connective for cases involving the perfect/non-perfect distinction is similar to the process described here for the cases involving the progressive/simple distinction. Thus, we take this more focused discussion to be representative of the general selection process and we will return to the perfect/non-perfect distinction shortly in Section 5.4.

Each temporal connecting word may correspond to several temporal interval relationships. Conversely, each temporal interval relationship corresponds to multiple temporal connecting words. For example, in terms of Allen’s temporal relations, the word *while* can represent =, *oi*, *s*, *d*, or *f*, and the temporal interval relationship *f* can be expressed as *after* or *while*. In addition, the non-inherent aspectual features of the matrix and adjunct verb can alter the meaning of the connecting word. For example, the progressive perspective of the verb endows the connecting word *before* with the possible meanings <, *o*, and *fi*. In the following sentences, *before* covers all three temporal interval relationships simultaneously:¹⁴

- (20) (i) Mary was drawing a circle before John was writing. (Dp/Dp)
- (ii) Mary was drawing a circle before John was sick. (Dp/Ss)
- (iii) John was sick before Mary was drawing a circle. (Ss/Dp)
- (iv) John was sick before Mary was unhappy. (Ss/Ss)

Since the matrix phrase is either a progressive event or a simple state, the adjunct phrase might start after the matrix finishes (<) or before the matrix finishes. If the adjunct phrase starts before the matrix finishes, it might finish at the same moment as the matrix (fi) or after the matrix (o). The interpretation changes significantly if the adjunct clause is realized in the simple past, in which case only the (<) reading is available:

- (21) (i) Mary was drawing a circle before John wrote a letter. (Dp/Ds)
- (ii) Mary was sick before John wrote a letter. (Ss/Ds)

¹³ Following Dowty (1979), we assume Sp does not exist since stative verbs are not realized in the progressive aspect.

¹⁴ Although the auxiliary *be* is used in the “drawing” and “writing” clauses of (20), we view the matrix verb to be non-stative. Inherent aspectual features are based on information associated with underlying lexical items, not on surface forms that result from their combination with other lexical items. This does not preclude the possibility that *non-inherent* aspectual features (e.g., coerced or non-lexical features) might be derived from the combination of lexical items (as described in the work of Verkuyl (1972), Olsen (1997), Dorr (1992), among others).

We have determined the possible temporal interval meanings associated with the $[\pm\text{dynamic}]/[\pm\text{progressive}]$ feature combinations through an extensive analysis of sample sentences such as (20)(i)–(iv) and (21)(i)–(ii). From this information, we have constructed *analysis charts*, which associate temporal interval meanings with connecting words for each $[\pm\text{dynamic}]/[\pm\text{progressive}]$ combination. The analysis chart for three representative temporal connectives (*after*, *before*, and *while*) is given for the Past/Past tense combination in Appendix B. These charts include more fine-grained aspectual categories, i.e., events, states, and processes, following the scheme suggested in column 3 of Table 1.

Because the distinction between events and processes has no impact on the choice of temporal connective, the analysis charts have been compiled into more succinct, two-dimensional *selection tables*—one for each connecting word. The selection tables for *after*, *before*, and *while* that apply to the Past/Past tense pairs are given in Table 5. For completeness, we also show the Past/Past selection table for *since*, *until*, and *when* in Table 6, although most of our examples focus on the selection of *after*, *before*, and *while*.

Each *row* of the selection table corresponds to a particular aspectual type/perspective and each *column* corresponds to a particular temporal interval. A ‘Y’ (= *yes*) signifies that the temporal connective covers the temporal interval meaning for that column given the pair of aspect values for that row. Since tense can also have an effect on the meaning of connecting words, a set of analysis charts must be constructed for each allowable matrix/adjunct tense pair. Appendix C provides an example of one additional analysis chart for *after*, *before*, and *while* and the corresponding (more succinct) selection tables for the future/present tense combinations.

Note that, if we include the perfect/non-perfect distinction in these selection tables, we have a more expanded analysis which includes two more combinations: Df and Sf (where f=perfect). The selection process uses the same table lookup approach for the perfect/non-perfect distinction as it does for the progressive/non-progressive distinction, as we will see shortly in Section 5.4.

Steps e–g of the algorithm in Figure 2 operate as follows. Step e selects an interval relation that gets used, along with values for $[\pm\text{dynamic}]$ and $[\pm\text{progressive}]$, to select a temporal connective. In step f, each selection table is inspected to determine whether its connecting word can be used. The selection tables are searched according to the interval relation, in order, from sparsest to densest. A connecting word with a sparse table has a more specific meaning than one with a dense table, since it has a narrow range of possible meanings. For example, consider an input with the following specification: Ss matrix, Ss adjunct, and temporal interval *f* (finishes). Although both the *after* and *while* tables contain a Y in the relevant cell (i.e., the one marked as row “Ss/Ss” and column “f”), *after* will be preferred over *while* since the associated table is the sparser of the two.

The use of the tables from sparsest to densest is an application of a “simplicity heuristic” that is used to rank the possible surface sentences. It is expected that a full text planning (e.g., (Hovy, 1993)) would replace this heuristic with a

Table 5. Selection Tables for Past/Past Tense Combination: AFTER, BEFORE, WHILE

AFTER													
	=	o	oi	s	si	d	di	m	mi	f	fi	<	>
Dp/Dp			Y							Y			Y
Dp/Ds													Y
Dp/Ss			Y							Y			Y
Ds/Dp			Y							Y			Y
Ds/Ds													Y
Ds/Ss													Y
Ss/Dp			Y							Y			Y
Ss/Ds													Y
Ss/Ss			Y							Y			Y

BEFORE													
	=	o	oi	s	si	d	di	m	mi	f	fi	<	>
Dp/Dp		Y									Y	Y	
Dp/Ds												Y	
Dp/Ss		Y									Y	Y	
Ds/Dp												Y	
Ds/Ds												Y	
Ds/Ss												Y	
Ss/Dp		Y									Y	Y	
Ss/Ds												Y	
Ss/Ss		Y									Y	Y	

WHILE													
	=	o	oi	s	si	d	di	m	mi	f	fi	<	>
Dp/Dp	Y		Y	Y		Y				Y			
Dp/Ds	Y		Y	Y		Y				Y			
Dp/Ss	Y		Y	Y		Y				Y			
Ds/Dp	Y		Y	Y		Y				Y			
Ds/Ds	Y		Y	Y		Y				Y			
Ds/Ss	Y		Y	Y		Y				Y			
Ss/Dp	Y		Y	Y		Y				Y			
Ss/Ds	Y		Y	Y		Y				Y			
Ss/Ss	Y		Y	Y		Y				Y			

Table 6. Selection Tables for Past/Past Tense Combination: SINCE, UNTIL, WHEN

SINCE												
	=	o	oi	s	si	d	di	m	mi	f	fi	<>
Dp/Dp									Y	Y		Y
Dp/Ds												
Dp/Ss												
Ds/Dp												
Ds/Ds									Y	Y		Y
Ds/Ss									Y	Y		Y
Ss/Dp												
Ss/Ds									Y	Y		Y
Ss/Ss									Y	Y		Y

UNTIL												
	=	o	oi	s	si	d	di	m	mi	f	fi	<>
Dp/Dp				Y				Y				
Dp/Ds				Y				Y				
Dp/Ss				Y				Y				
Ds/Dp				Y				Y				
Ds/Ds				Y				Y				
Ds/Ss				Y				Y				
Ss/Dp				Y				Y				
Ss/Ds				Y				Y				
Ss/Ss				Y				Y				

WHEN												
	=	o	oi	s	si	d	di	m	mi	f	fi	<>
Dp/Dp	Y	Y			Y		Y				Y	
Dp/Ds	Y	Y		Y	Y		Y	Y			Y	Y
Dp/Ss	Y	Y		Y	Y		Y					
Ds/Dp	Y		Y			Y			Y	Y		
Ds/Ds	Y		Y	Y	Y	Y			Y	Y		Y
Ds/Ss	Y		Y	Y	Y	Y			Y	Y		
Ss/Dp	Y		Y	Y	Y	Y			Y	Y		
Ss/Ds	Y		Y	Y	Y	Y			Y	Y		Y
Ss/Ss	Y		Y	Y	Y	Y			Y	Y		

more pragmatic selection technique. In a full generation model (which includes a discourse component), the speaker might want to convey a causality component of meaning, possibly reordering clauses and using different connectives. For example, the speaker might want to reorder the clauses in *John was sick before Mary was happy* and use the connective *while* or *when* instead of *before*: *Mary was unhappy while/when John was sick*. Since our goal is to produce legal surface combinations, without knowledge of context or discourse rules, we take the simplicity heuristic to be a reasonable approximation to an appropriate solution without the drawback of additional complexity. It would be overkill to apply techniques for ruling out sentences that are legal at this stage of the process, given that CONGEN is intended to be used as a sub-component of a full MT system where overgeneration is constrained by linguistically motivated rules and statistical extraction techniques (Habash and Dorr, 2002).

Once step f in Figure 2 has selected a ranked set of appropriate temporal connectives for each BTS pair, step g applies a corpus-based ranking of each connective set. Section 7 describes the corpus-based ranking of temporal connectives in more detail.

5.4 Application of Algorithm to Ongoing Example

We shall complete the application of the algorithm in Figure 2 to our example:

(22) $\text{go}(\text{john}, \text{store}, 15:00, 15:15) \wedge \text{arrive}(\text{mary}, 15:31, 15:31)$

To recap, we have already determined that the ranked list of BTS pairs P is $\{(\text{past}, \text{past}), (\text{past}, \text{past perfect}), (\text{past perfect}, \text{past}), (\text{present perfect}, \text{present perfect}), (\text{past perfect}, \text{past perfect})\}$, the aspectual perspective pairs A is $\{(\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), \}$, and the temporal relationship T between the literals is *before* ($<$). The sentences corresponding to these settings were given in (18), repeated here for convenience:

- (23) (i) John went to the store *CW* Mary arrived.
(ii) John went to the store *CW* Mary had arrived.
(iii) John had gone to the store *CW* Mary arrived.
(iv) John has gone to the store *CW* Mary has arrived.
(v) John had gone to the store *CW* Mary had arrived.

Now steps f–g must determine the appropriate temporal connectives for each of the possibilities in P, given the temporal relationship $<$. For example, case (23)(i) corresponds to the Ds/Ds combination with a (past, past) tense assignment; this means that step f must examine all (past, past) selection tables until it finds a “Y” the cell associated with the Ds/Ds row and the “ $<$ ” column in the past/past table (see Tables 5 and 6). The only applicable connective is *before*, which means *CW* may be replaced by *before* in (23)(i).

The next four cases involve the perfect aspect, which require access to new information not previously presented in our selection tables. Given that it is possible to generate a non-perfect form of the sentence, it is clear that the other

four cases will have a lower ranking. However, for the purpose of illustration, we show the additional rows of the expanded *before* tables for Past/Past and Pres/Pres in Table 7—where Df is used for [+dynamic]/[+perfect].¹⁵ Note, in particular, that sentence (23)(iv) (the only bad one) is ruled out as a candidate for *before* because the Pres/Pres selection table disallows the Df/Df combination under the “<” column. On the other hand, if Df/Ss had been selected for the clauses in our example, this combination would have been allowed by the Pres/Pres *before* table: *John has gone before Mary arrives*.

Table 7. Selection Tables Containing Perfect Aspect: BEFORE

BEFORE: Past/Past											
	=	o	oi	s	si	d	di	m	mi	f	fi
Df/Df											Y
Df/Dp											
Df/Ds											Y
Df/Ss											Y
Dp/Df											
Ds/Df											Y
Ss/Df											

BEFORE: Pres/Pres											
	=	o	oi	s	si	d	di	m	mi	f	fi
Df/Df											
Df/Dp											
Df/Ds											
Df/Ss											Y
Dp/Df											
Ds/Df											
Ss/Df											

Examples: *John went to the store before Mary had arrived* (Ds/Df); *John had gone to the store before Mary arrived* (Df/Ds); * *John has gone to the store before Mary has arrived* (Df/Df). *John had gone to the store before Mary had arrived* (Df/Df);

The omission of Df/Df in the Pres/Pres table for *before* might be seen as problematic since native English speakers consider sentences such as *John has gone to the store before Mary has arrived* to be acceptable. However, if this sentence is legal, it cannot apply to a situation in which both the matrix and adjunct clause occur in the past (i.e., prior to speech time) as in the current example. This might indicate that the status of $E_{sf_R,S}$ (present perfect) as an allowable BTS for the temporal relationship $E_s < E_f < S$ (see Table 3) is questionable. We leave this as a question for future investigation. In the meantime,

¹⁵ More complex combinations such as co-occurrence of Df and Dp (e.g., *John has gone to the store before Mary has been arriving*) are handled simply through access to multiple rows in the table, thus providing a powerful mechanism for constraint application.

our corpus-based experiments indicates that the omission of Df/Df is justified in the Pres/Pres *before* table, given that this combination did not arise.

It is important to note that this same point does not necessarily apply to the past perfect, which clearly adheres to the relationship $E_s < E_f < S$. Unlike the present perfect, the past perfect has an important role when considered in the context of *when* clauses, e.g., *John had gone to the store when Mary arrived* vs. *John went to the store when Mary arrived*. We will see in Section 7 that such combinations are prevalent in our corpus analysis; the expanded Past/Past selection table for *when* covers such cases. The additional rows are shown in Table 8.

Table 8. Past/Past Selection Table Containing Perfect Aspect: WHEN

WHEN: Past/Past												
	=	o	oi	s	si	d	di	m	mi	f	<	>
Df/Df	Y		Y	Y	Y		Y		Y	Y		Y
Df/Dp	Y		Y	Y	Y		Y		Y	Y		Y
Df/Ds	Y		Y	Y	Y		Y		Y	Y		Y
Df/Ss	Y		Y	Y	Y		Y		Y	Y		Y
Dp/Df	Y		Y	Y	Y		Y		Y	Y		Y
Ds/Df	Y		Y	Y	Y		Y		Y	Y		Y
Ss/Df	Y		Y	Y	Y		Y		Y	Y		Y

Given the expanded selection tables, Step f produces the list C of singleton sets of connectives for each BTS: {(before), (before), (before), (before)}. Step g applies trivially to each of these sets. Thus, the algorithm in Figure 2 outputs $P = \{(\text{past}, \text{past}), (\text{past perfect}, \text{past}), (\text{past perfect}, \text{past perfect}), (\text{past}, \text{past perfect})\}$; $A = \{(\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), (\text{simple}, \text{simple}), \}$, and $C = \{(\text{before}), (\text{before}), (\text{before}), (\text{before})\}$. These outputs correspond to the following ranked list of sentences:

- (24) (i) John went to the store before Mary arrived.
(ii) John went to the store before Mary had arrived.
(iii) John had gone to the store before Mary arrived.
(iv) John had gone to the store before Mary had arrived.

It should be noted that, like the sparsest-to-densest approach, the preference for *non-perfect* over *perfect* could be replaced by a more pragmatic selection technique. While the use of such techniques are outside of the scope of this paper, we have observed a correspondence between our preference scheme and usage in a comprehensive corpus analysis. (This point is discussed further in Section 7.)

6 Extended Examples

This section examines three additional examples of the process for selecting tense, aspect, and connecting words during sentence generation.

Example 1 Consider the following conjunction of two time-stamped literals.

(25) $\text{curious}(\text{Mary}, t_3, t_4) \wedge \text{hide}(\text{John}, \text{book}, t_1, t_2)$

where t_1 is later than S and $t_1 < t_2 < t_3 < t_4$. We assume that $\text{curious}(\text{Mary}, t_3, t_4)$ will be mapped into the matrix clause and $\text{hide}(\text{John}, \text{book}, t_1, t_2)$ will be mapped into the adjunct clause.¹⁶

Each literal is associated with a closed interval that occurs after the speech time, so according to the mapping in Table 3 (step a of Figure 2), each literal has a set of three BTSs: {future, future perfect, present}. Using the CDTS in Table 4 (step b of Figure 2), an allowable adjunct tense for each basic tense is determined: {(future, present), (future, future), (future, future perfect), (future perfect, present), (future perfect, future), (future perfect, future perfect), (present, present)}. Step c ranks these choices using our corpus-based analysis (to be described in the next section). The result is the following (re-ordered) list: {(present, present), (future, present), (future, future), (future perfect, present), (future, future perfect), (future perfect, future), (future perfect, future perfect)}. (The last three cases are ranked equivalently, as these combinations did not occur at all in our corpus.)

Thus, given sufficient grammatical information about the two literals, the following sentences could be generated (ignoring aspect selection) to fit the matrix/adjunct tense pairs:¹⁷

- (26) (i) Mary is curious *CW* John hides / is hiding the book.
(ii) Mary will be curious *CW* John hides / is hiding the book.
(iii) Mary will be curious *CW* John will hide.
(iv) Mary will have been curious *CW* John hides / is hiding the book.
(v) Mary will be curious *CW* John will have hidden the book.
(vi) Mary will have been curious *CW* John will have hidden the book.
(vii) Mary will have been curious *CW* John will have hidden the book.

Let us consider the selection of progressive vs. simple for this conjunction of literals. The first literal corresponds to the [-dynamic] verb *be curious* and must be realized in the simple form, according to constraint I of Figure 4.¹⁸ As for the second literal, the corresponding verb *hide* is inherently [-atomic]. Given that the event occurs in a closed non-point interval, constraint IV (a) of Figure 4

¹⁶ The alternative ordering can be tried if no connective possibilities work out.

¹⁷ The phrase *be curious* is viewed as a single stative verb corresponding to the predicate *curious*. For simplification, we ignore the possibility of using progressive for this verb since it will later be ruled out by constraint I in Figure 4.

¹⁸ Since *simple* refers to non-progressive, the following cases are all considered simple: *is curious*, *will be curious*, and *will have been curious*.

applies and *hide* is realized in the simple form—thus eliminating the progressive forms shown above in (26).

Now there are seven possible surface sentences and the final step is to choose a temporal connective. According to the interval relationships in Table 2, the two literals are in a “>” (after) relation. Note that sentence (26)(ii) is a Ss/Ds pair occurring as a Future/Present combination. The selection table for the Future/Present combination (at the end of Appendix C) contains a “Y” entry under the “>” column only in the case of the connective *after*. Thus, this sentence is realized with *after* as the temporal connective:

(27) Mary will be curious after John hides.

The connective *after* is also available for the other combinations—except for case (i), which is thus eliminated in step f. However, just as the simple perspective is generally ranked higher than the perfect perspective, the non-complex forms (e.g., *hides* in (ii)) are ranked higher than complex phrases (e.g., *will hide* or *will have hidden* in (iii)–(vii)). Although this choice appears to be unprincipled, in fact it is a close approximation to what we observed in our corpus analysis, as we will see in Section 7.¹⁹ Thus, once it is confirmed that a temporal connective is available for the less complex form in (27) above, this is the one that is ranked highest. The overall ranking of sentences is:

- (28) (i) Mary will be curious after John hides
- (ii) Mary will be curious after John will hide.
- (iii) Mary will have been curious after John hides
- (iv) Mary will be curious after John will have hidden the book.
- (v) Mary will have been curious after John will have hidden the book.
- (vi) Mary will have been curious after John will have hidden the book.

Example 2 Now let us assume the literal *hide(john,book,t₁,t₂)* underlies the matrix clause. The tense and aspectual features remain unchanged, but they will be reversed (i.e., Ds/Ss), so the temporal relationship is <, and the only possible connecting word is *before*:

(29) John will hide the book before Mary is curious.

Example 3 Now let us consider the same case again (with *hide(john, book, t₁,t₂)* as the matrix) but with the following start/end relationships: $t_1 < t_3$, $t_4 < t_2$, and $t_2 < S$. Again, the aspect values remain simple for both literals and

¹⁹ The use of future tense in both clauses is, in any case, problematic. As noted in (Hornstein, 1990, p. 214, fn.23), “It has long been observed that the future is not felicitous in adverbial adjuncts.” (This is also confirmed in our corpus analysis.) Researchers like Yip have suggested a “principle of economy” that deletes *will* in order to account for such cases but, as Hornstein argues, such a solution is not systematic since it does not apply to other tenses. Our own approach is to produce both cases, but to eliminate the non-economical one if the temporal connective allows a more economical solution.

the matrix/adjunct designation is still Ds/Ss. But now the temporal relationship between the matrix and adjunct clauses is *di*. No connecting word satisfies Ds/Ss and *di* in the Future/Present combination (at the end of Appendix C). However, if the *curious* literal is chosen as the matrix, the Ss/Ds and *d* combination specifies *while*. The resulting sentence for the conjunction is:

(30) Mary will be curious while John hides the book.

7 Corpus-Based Verification of Theoretical Results

We conducted a corpus-based analysis in order to verify both Hornstein’s analysis and our own approach to generating temporally conjoined clauses. Our approach to theoretical verification involved several “counting” experiments. We built a Perl script to analyze all 53,411 sentences in the tagged Lancaster-Oslo-Bergen (LOB) corpus.²⁰ Of these, 3592 contained Matrix/Adjunct pairs conjoined by a temporal connective, i.e., a total of 7184 clauses.²¹

First we examined tense combinations in the LOB. The tenses of each clause were easily detectable by means of the tags shown in Table 12. Of the 7184 clauses, 3878 were unambiguously tensed. The number of occurrences of simple/perfect/progressive for each of these 3878 tenses is shown in Table 9. In this table, the numbers for Matrix and adjunct clauses are separated, e.g., the number 40/28 refers to 40 Matrix and 28 Adjunct which contain the simple/future form. Note that the simple tenses were 8-10 times more likely to occur than the progressive or perfect. This difference in usage strongly supports a preference for simple over the progressive or perfect. Note also that the LOB sentences include far fewer future-tense clauses than past- or present-tense clauses. This is most likely a corpus-specific finding, thus suggesting that our approach may be further enhanced by fine-tuning the tense preferences according to the corpus style.

Table 9. Matrix/Adjunct Tense Combinations in LOB

Tense/Perspective	Simple	Perfect	Progressive
Future	40/28	1/2	2/2
Past	864/1231	91/166	47/56
Present	462/679	70/77	14/46

Next we verified the CDTS in our corpus analysis. Of the 3592 Matrix/Adjunct pairs, only 1033 of these were unambiguously tensed in *both* the Matrix and Adjunct clauses.²² We counted the number of occurrences of different types of

²⁰ The code used to process the corpus is available from either of the two authors.

²¹ All results of our corpus-based analysis are stored in <http://www.umi.acs.umd.edu/~bonnie/corpus-results.txt>.

²² Also, included in the 3592 Matrix/Adjunct pairs were some modal clauses, such as *should go*, which we are not including in this analysis. Modals are, indeed, a part of

Matrix/Adjunct clauses and checked for consistency with Hornstein’s analysis. The results are grouped into *disallowed* (CDTS violation) and *allowed* (no CDTS violation) as shown in Table 10. The *allowed* cases are used for our corpus-based ranking of BTS pairs in step c of Figure 2.

Table 10. Disallowed and Allowed Tense Combinations in LOB

Disallowed		Allowed	
Tense Pair	Count	Tense Pair	Count
Future/Past	0	Future/Present	24
Future/Past Perfect	0	Future/Present Perfect	5
Future Perfect/Past	0	Future/Future	2
Future Perfect/Past Perfect	0	Future/Future Perfect	0
Past/Future	0	Future Perfect/Present	1
Past/Future Perfect	0	Future Perfect/Present Perfect	0
Past/Present	15	Future Perfect/Future	0
Past/Present Perfect	0	Future Perfect/Future Perfect	0
Past Perfect/Present	3	Past/Past	571
Past Perfect/Present Perfect	0	Past/Past Perfect	51
Past Perfect/Future	0	Past Perfect/Past	21
Past Perfect/Future Perfect	0	Past Perfect/Past Perfect	1
Present/Future	4	Present/Present	237
Present/Future Perfect	0	Present/Present Perfect	19
Present/Past	29	Present Perfect/Present	11
Present/Past Perfect	4	Present Perfect/Present Perfect	12
Present Perfect/Past	21		
Present Perfect/Past Perfect	1		
Present Perfect/Future	1		
Present Perfect/Future Perfect	0		
TOTAL	78	TOTAL	955

The results of this experiment indicate that disallowed matrix/adjunct pairs occurred rarely (an average of 4 occurrences per pair) in comparison to allowed matrix/adjunct pairs (an average of 60 occurrences per pair). Thus, our analysis provides strong support for our use of Hornstein’s CDTS in generating surface sentences. We see this as a success: The cases not accounted for in our theoretical framework amount to a mere 7% of the sentences studied, well within the range of broadscale applicability of the approach.

It is worthwhile to examine the disallowed cases, so that we have a better understanding of where (and why) the current approach would fall short. Only three cases occurred with higher-than-average frequency: Past/Present (15 times), Present/Past (29 times), and Present Perfect/Past (21 times). Ap-

Hornstein’s analysis (pp. 33–38) but the tenses of these were not as easily detectable from the corpus, so we ran the experiment on the non-modals.

pendix D examines each of these cases, in turn, and provides detailed examples of each case.

Finally, we did an analysis of the temporal connectives (such as *since*) in the LOB. These were marked with the special tag “CS”. The most frequently occurring temporal connectives from the corpus were *after*, *before*, *since*, *until*, *when*, and *while*. We hand-categorized each of these connectives in terms of Allen’s temporal relations, based on a subset of the sample sentences in LOB. The connectives are shown in Table 11 with their frequency of occurrence and possible temporal relations.²³

Table 11. Temporal Connectives: Frequency of Occurrence (LOB) and Allen’s Interval Relations

Connective	Frequency	Interval Relations
after	143	>,oi,f
before	354	<,o,fi
since	264	>,mi,f
until	269	m,s
when	2118	=,o,oi,s,si,d,di,m,mi,f,fi,>
while	444	=,o,oi,s,si,d,f

It is important to note that the distinction between temporal connectives and causal connectives falls out quite naturally from our approach. In particular, the connective *when* often functions as a causal connective rather than a temporal connective, e.g., *John laughed when Mary fell*. As discussed previously by Moens and Steedman (1988), *when* has a strong causality component of meaning; this is clearly beyond the scope of this work. The use of connecting words that convey causal or other meanings may introduce ambiguity into a generated sentence.²⁴ The elimination of such ambiguity is an important open issue, but our current approach addresses this naturally, in that the simplicity heuristic generally ranks less frequently occurring (less ambiguous) connectives (e.g., *after*) higher than frequently occurring (more ambiguous) connectives (e.g., *when*).

In addition to frequency counting of the 3592 Matrix/Adjunct corpus pairs, we automatically categorized each pair, along with its mediating connective, into

²³ Other connectives that we studied, but did not include in our corpus analysis, were taken from Webster’s 7th Collegiate Dictionary (Merriam-Webster, 1976)). These included: *as long as* (=,f); *as of* (s); *as soon as* (>,mi); *during* (d,s,f); *ere* (<); *following* (>); *now that* (>,mi); *once* (>,mi); *past* (>,si); *prior* (m,<); *pursuant* (mi); *so long as* (=,f); *till* (m,s).

²⁴ For example, it is difficult to tell which meaning is assigned to the word *when* in *John laughed when Mary fell*. If the causal meaning is intended, then it is likely that the laughing occurred only *after* the fall, unlike the temporal reading, where the laughing is interpreted to occur during the time of the fall as well as shortly thereafter.

Table 12. Tense Tags in LOB

Tense	Tags	Example
Pres	VB	give
Pres Perf	HV VBN	have urged
Past	VBD	started
Past Perf	HVD VBN	had stated
Fut	FUT VB	will see
Fut Perf	FUT HV VBN	will have appeared

one of 144 cases so that we could verify our hand-generated connective-selection tables. The 144 cases were derived from 12 tense specifications (Df, Dp, Ds, and Ss—multiplied out 3 times for each tense), each one crossed with the same 12 tense specifications. The tense specifications were detectable by the tags given above in Table 12 together with additional tags for distinguishing Stative (Ss) from Dynamic (Df, Dp, Ds). (A clause was given the status Ss when its main verb was labeled BEZ, e.g., *is late* or BEDZ, e.g., *was late*; it was otherwise given the status of Df, Dp, or Ds.) A frequency count was stored in each cell of each table. A very small subset of the 144 pairs (just the Past/Past combinations) is shown in Table 13 for each of the six connectives in our study. The fully expanded version of this table (with 144 entries for each connective) is used for our corpus-based ranking of temporal connectives in step g of Figure 2.

These frequency-count tables were compared against our human-generated selection tables (e.g., Tables 5 and 6 and the Past/Past tables containing perfect aspect 7 and 8). We found that our tables account for many cases that do not show up in the corpus, there were no corpus cases examples that could not be account for by our selection tables. For example, the selection table for *since* (see Table 6) excludes the Dp/Dp, Dp/Ss, and Ds/Dp combinations; these were precisely the combinations that never arose in our corpus analysis of *since*. On the other hand, the Past/Past selection table for *after* (Table 5) includes the Ds/Dp combination, even though this did not arise in our corpus, because sentences like *Mary drew a circle after John was writing a letter* passed human inspection prior to our corpus investigation.

It should be noted, however, that many cases in the human-generated selection tables that did not arise in the corpus are in the category of “borderline acceptable,” e.g., the Dp/Ss combination for *after*: *John was angry after Mary was drawing a circle*. It is clear from such cases that the corpus-based decision made in step g of Figure 2 is an important step toward producing a reasonable ranking of output possibilities for the final selection. This suggests that more improved results might be obtained if the selection tables were built automatically in the first place, using corpus-based techniques. Taking this approach would allow for domain-specific tuning so that output options would more closely match the contents of the particular corpus that is selected.

Table 13. Corpus-Induced Frequency Counts for 16 Past/Past Combinations Using *after*, *before*, *since*, *until*, *when*, and *while*

Matrix/ Adjunct	AFTER	BEFORE	SINCE	UNTIL	WHEN	WHILE
Df/Df	1	2	3	2	1	1
Df/Dp					1	
Df/Ds		6	3	5	17	
Df/Ss		1			13	2
Dp/Df						
Dp/Dp						
Dp/Ds			2	1	24	1
Dp/Ss		3		2		1
Ds/Df	15	3	1	4	12	
Ds/Dp				3	5	7
Ds/Ds	8	25	2	30	168	31
Ds/Ss	4	6	1	12	45	13
Ss/Df	1		5	1	4	1
Ss/Dp					2	
Ss/Ds	2	18	2	1	90	8
Ss/Ss		6	2	1	36	7

8 Discussion

The processes described above for selecting tense, aspect and connecting words build on the theoretical work of Allen and Hornstein. It is useful to consider what we would, and would not, have been able to achieve if we had used only one theory or the other in the work.

The two theories are complementary. From Hornstein’s work we gain a careful analysis of tenses and how they fit together in complex sentences. From Allen’s work we gain the ability to represent and manipulate time information at an abstract level separate from individual time-stamps. Also, we gain the ability to represent events that have duration.

If we were to omit Allen’s theory, we would lose the ability to deal with information about intervals. All events would have to be considered as point events. The only relevant temporal connecting words would be those that express $t_1 < t_2$ (*before*), $t_1 = t_2$ (*at the same time as*), and $t_1 > t_2$ (*after*), where t_1 and t_2 are the timepoints of the events. Furthermore, we would have no means to select values for non-inherent aspectual feature values.

Alternatively, if we were to omit Hornstein’s theory, we would have no way to generate tense structures for matrix/adjunct sentences in a general way. When temporal connectives are considered in languages other than English, more evidence emerges to support the claim that inherent aspectual values have an effect on the meanings of the temporal connectives (Dorr, 1992). Although we could analyze complex tenses on a case by case basis using Allen’s work, we would

have to redo the analysis for languages other than English. Verification of the analysis would be difficult.

By amalgamating Allen’s theory of intervals with Hornstein’s theory of tense, we arrive at a principled theory for selecting basic tenses for events that take place over intervals as well as points. We use Allen’s representation to extend Hornstein’s tense analysis for the purpose of selecting tenses for events during the generation of language. From this, we have developed a comprehensive method to select basic tenses for durative events as well as instantaneous events. The two theories together enable us to select aspectual feature values for events based on their temporal intervals and on the tense selected for the event. Allen’s theory provides a means to semantically analyze and select a wider range of temporal connectives than if we used only the three relationships mentioned above that apply to point events.

Our tense selection process is a three step procedure. Two tables have been compiled from Hornstein’s theory, thus reducing the first two steps of the tense selection process to a simple table look-up procedure. The selection of temporal connecting words is based on a set of tables, one for each connecting word, that have been compiled through a careful process of human analysis based on sentences such as those shown in Appendices B and C. The meaning of temporal connectives may change in different aspectual contexts, and this information is available in the tables.

Beyond Allen’s and Hornstein’s theories, the aspect selection process relies on our distinction between *inherent* and *non-inherent* aspectual features of verbs. Inherent features are important for the selection of a particular verb token for the surface sentence. They are also important in that they can implicitly represent temporal information such as duration. If a verb conveys certain temporal information implicitly, then a generated sentence may not need to include the temporal information explicitly. Non-inherent feature values may be selected during generation, depending on inherent features and on temporal information associated with the literals. A supporting source of information for more refined lexical choice is described in (Dorr and Voss, 1996; Dorr and Olsen, 1996), where verb choices are narrowed down according to the associated aspectual information.

9 Conclusions and Future Work

This paper describes the selection of tense, aspect, and connecting words using a general method for processing temporal information in the generation of language. The ability to handle time is not only essential to interfaces for accessing temporal databases or document collections (as in Q/A), but it is also essential in other applications such as MT since language cannot be produced without tense and aspect assignment.

Our approach is limited to the generation of single matrix/adjunct sentences that describe two events. To extend the approach to multiple events, two directions must be pursued: (1) the connection of three or more events in a single

sentence, and (2) the discourse level planning of which events to connect, which events should stand alone, and what order to put them in. In addition, the study of other categories of connecting words such as causal and spatial connectives (Elhadad and McKeown, 1990) could further enhance a multi-event approach.

Although our focus is on *generation*, our work could benefit from additional temporal processing techniques developed by researchers who are addressing *analysis* of temporal expressions, e.g., the framework of (Androutsopoulos, 1999), where a representation language called TOP is used to provide the semantics behind entries in temporal databases; the techniques developed by Crouch and Pulman (1993) and Hwang and Schubert (1994) for interpreting temporal information in natural language expressions; the work of Lascarides and Oberlander (1993), where temporal connectives are analyzed in a discourse context; the automatic semantic-tagging approach of Schilder (1999) and Schilder and Habel (2001), where presuppositions are derived from temporal connectives; and the model of Steedman (1997), where compositional aspectual knowledge is used for automatic temporal reference resolution.

The main results of this paper are the following. We have provided a theory for selecting tenses for individual events that may be either points or intervals in time. The selection theory extends the framework of Hornstein (1990) using the temporal interval representation of Allen (1983) and (Allen, 1984). For literals that are to be combined in a matrix/adjunct structure, selected tenses are constrained by Hornstein’s constraint on derived tense structure. Next, we have provided a theory for aspect selection that is constrained by the tenses already selected for an event; the aspectual constraints are adapted from the work of Dowty (1979). Finally, we have provided a methodology for selecting connecting words through access to a set of tables that associate temporal interval relationships with combinations of connecting word and aspectual values. The connecting word selection is constrained by the aspectual values already selected for an event and also on preferences induced from an extensive corpus analysis.

The theoretical results described here serve as the basis of a temporal framework for an implemented system that generates English sentences from time-stamped literals (Dorr, Habash, and Traum, 1998). Our future work will entail the complete integration of CONGEN into the GHMT framework (Habash and Dorr, 2002) and the use of more comprehensive tense-selection techniques (Olsen et al., 2000; Olsen et al., 2001). However, more important than providing a critical component for generation of temporal expressions is the understanding that we have gained about interconnections between tense, aspect, and temporal connecting words. We believe this understanding provides an adequate basis for the development of additional processing modules and for the extension of linguistic coverage in the field of language generation.

Acknowledgments

Bonnie J. Dorr’s participation in this project was supported in part by NSF grants IRI-9357731/9629108 (NYI/PFF), Alfred P. Sloan Fellowship BR3336, a DoD Contract, and DARPA N66001-97-C-8540/N66001-00-2-8910. Terry Gaasterland’s participation in this project was supported in part by the DOE Office of Scientific Computing, contract W-31-109-Eng-38. We would like to thank Julia Koslovsky for her assistance in building and running a Perl script to analyze all 53K+ sentences in the tagged Lancaster-Oslo-Bergen (LOB) corpus.

A Basic Tense Structures for Events with Duration

To extend the BTS framework to cover events that have duration over some interval in time, we consider the event E to have two time-stamps, a starting time-stamp E_s and a finishing time-stamp E_f . Either of E_s and E_f may be *open* or *closed*. We assume speech time S to be a point interval, i.e., an interval with identical start and stop times. We shall now inspect each configuration of E_s and E_f and S for E_s and E_f open and closed, respectively.

Consider the ordering $E_s < E_f < S$. The interval relationship is $E_{sf} < S$ as in:

$$(31) \quad \begin{array}{ccc} E_s & E_f & S \\ \bullet & \text{---} \bullet & \bullet \end{array}$$

For this configuration, three BTSs preserve the relationship: past E, R, S , past perfect E, R, S , and present perfect E, S, R .

For the opposite configuration,

$$(32) \quad \begin{array}{ccc} S & E_s & E_f \\ \bullet & \text{---} \bullet & \bullet \end{array}$$

there are three possible tenses that preserve the relationship between E and S : future S, R, E , future perfect S, E, R , and present S, R, E .

When the start time for a literal precedes the speech time and the stop time follows the speech time, as in the following:

$$(33) \quad \begin{array}{ccc} E_s & S & E_f \\ \bullet & \text{---} \bullet & \bullet \end{array}$$

a decision can be made to focus on either the starting point or the finishing point, in which case that point is used to determine the possible basic tenses.

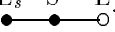
When either the start time or the stop time coincides with the speech time, as in the following:

$$(34) \quad \begin{array}{ccc} E_s & E_f \\ \bullet & \text{---} \bullet \\ S & \bullet \end{array}$$

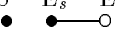
$$(35) \quad \begin{array}{ccc} E_s & E_f \\ \bullet & \text{---} \bullet \\ S & \bullet \end{array}$$

a decision must be made about whether to focus on the entire event’s relationship to S or on either start time or stop time.

Now let us consider intervals that have an open stop time. A literal with an open stop time describes an event or state that is ongoing. If the start time precedes speech time, as in the following:

$$(36) \quad E_s \quad S \quad E_f$$


the event or state is true now and will continue to be true for some period of time. A decision must be made whether to focus on the start time (in which case, E_s is used to determine the set of tenses) or to focus on the ongoing event (in which case, $E_{sf} \leq S$ is used). Alternatively, if the start time is in the future, the whole event is in the future.

$$(37) \quad S \quad E_s \quad E_f$$


Finally, let us consider intervals that are actually points: i.e., $E_s = E_f$. For $E_{sf} < S$, the set of past tenses can be used. For $E_{sf} = S$, only the present tense can be used. For $E_{sf} > S$, the set of future tenses can be used.

Table 3 in Section 4 shows the compilation the possible BTSs for closed intervals, points, and open intervals for each possible relationship between the interval and S .

B Construction of Analysis Charts for Connecting Words

Section 5.3 discussed selection tables that were constructed for the connecting words *after*, *before*, and *while*, for the past/past tense combination in English. This appendix contains the analysis charts that were used to produce these condensed selection tables. In our analysis, we included more fine-grained aspectual categories, i.e., events, states, and processes, following the scheme suggested in column 3 of Table 1, but we found no distinction between events and processes; thus, the selection tables are consistently more succinct.

The analysis charts contain sentences that combine the aspectual properties *dynamic/state*, *progressive/simple* and *process/event*. Since sentences that are both *stative* and *progressive* do not occur in English, there are only 5 possible combinations for a single concept: dynamic/event/progressive, dynamic/event/simple, dynamic/process/progressive, dynamic/process/simple, state/simple. We use the following verb phrases to construct the sentences:

to write a letter (Event)

to draw a circle (Event)

to walk (Process)

to laugh (Process)

to be angry (State)

to be happy (State) Each analysis chart contains 25 examples since the matrix and adjunct clauses each have 5 possible realizations.

The analysis chart for the Past/Past tense combination is given here:

Matrix	AFTER	BEFORE	WHILE	Adjunct
Mary was drawing a circle <i>Event, Prog: ●—●</i>	> f oi	< fi o	= d s	John was writing a letter <i>Event, Prog: ●—●</i>
Mary was drawing a circle <i>Event, Prog: ●—●</i>	"	"	"	John was laughing <i>Process, Prog: ●—●</i>
John was laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary was drawing a circle <i>Event, Prog: ●—●</i>
John was laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary was walking <i>Process, Prog: ●—●</i>
Mary was drawing a circle <i>Event, Prog: ●—●</i>	>	<	= oi s d f	John wrote a letter <i>Event, Simp: ●</i>
Mary was drawing a circle <i>Event, Prog: ●—●</i>	"	"	"	John laughed <i>Process, Simp: ●</i>
John was laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary drew a circle <i>Event, Simp: ●</i>
John was laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary walked <i>Process, Simp: ●</i>
Mary was drawing a circle <i>Event, Prog: ●—●</i>	> f oi	< fi o	s d f = oi	John was angry <i>State, Simp: ●—●</i>
Mary was laughing <i>Process, Prog: ●—●</i>	"	"	"	John was angry <i>State, Simp: ●—●</i>
Mary drew a circle <i>Event, Simp: ●</i>	> f oi	<	= oi s d f	John was writing a letter <i>Event, Prog: ●—●</i>
Mary drew a circle <i>Event, Simp: ●</i>	"	"	"	John was laughing <i>Process, Prog: ●—●</i>
John laughed <i>Process, Simp: ●</i>	"	"	"	Mary was drawing a circle <i>Event, Prog: ●—●</i>
John laughed <i>Process, Simp: ●</i>	"	"	"	Mary was walking <i>Process, Prog: ●—●</i>
Mary drew a circle <i>Event, Simp: ●</i>	>	<	= s d f	John wrote a letter <i>Event, Simp: ●</i>
Mary drew a circle <i>Event, Simp: ●</i>	"	"	"	John laughed <i>Process, Simp: ●</i>
John laughed <i>Process, Simp: ●</i>	"	"	"	Mary drew a circle <i>Event, Simp: ●</i>
John laughed <i>Process, Simp: ●</i>	"	"	"	Mary walked <i>Process, Simp: ●</i>
Mary drew a circle <i>Event, Simp: ●</i>	>	<	s d f =	John was angry <i>State, Simp: ●—●</i>
Mary laughed <i>Process, Simp: ●</i>	"	"	"	John was angry <i>State, Simp: ●—●</i>
John was angry <i>State, Simp: ●—●</i>	> f oi	< fi o	= oi s d f	Mary was drawing a circle <i>Event, Prog: ●—●</i>
John was angry <i>State, Simp: ●—●</i>	"	"	"	Mary was walking <i>Process, Prog: ●—●</i>
John was angry <i>State, Simp: ●—●</i>	>	<	= s d f	Mary drew a circle <i>Event, Simp: ●</i>
John was angry <i>State, Simp: ●—●</i>	"	"	"	Mary walked <i>Process, Simp: ●</i>
John was angry <i>State, Simp: ●—●</i>	> f oi	< fi o	= s d f	Mary was happy <i>State, Simp: ●—●</i>

C Analysis Chart and Selection Tables for Future/Present Tense Combination

To build a full implementation of the method for selecting tense, aspect, and connecting words, an analysis chart must be built for all legal tense combinations. We apply the same analysis to obtain the chart for a future tense matrix combined with a present tense adjunct, e.g., *She will draw a circle while John is sleeping*. This analysis chart is then converted into selection tables for *after*, *before*, and *while*. This section shows our analysis chart and selection tables for the future/present combination.

Note that, in producing the selection tables from the analysis chart, we were able to take advantage of certain linguistic generalizations. In particular, we observed that the *event/process* distinction did not affect the connecting word meanings. Thus, we were able to construct a more succinct two-dimensional selection table indexed on one dimension by *dynamic/state-progressive/simple* combinations and on the second dimension by temporal intervals. We use the following abbreviations:

Dp = dynamic progressive

Ds = dynamic simple

Ss = state simple

Given two events, each with values for *dynamic/state* and *progressive/simple*, and given a temporal interval relation between the two events, we use the selection table to find an appropriate connecting word.

The analysis chart for the Future/Present tense combination is given here:

Matrix	AFTER	BEFORE	WHILE	Adjunct
Mary will be drawing a circle <i>Event, Prog: ●—●</i>	> f oi	< fi o	= d s	John is writing a letter <i>Event, Prog: ●—●</i>
Mary will be drawing a circle <i>Event, Prog: ●—●</i>	"	"	"	John is laughing <i>Process, Prog: ●—●</i>
John will be laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary is drawing a circle <i>Event, Prog: ●—●</i>
John will be laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary is walking <i>Process, Prog: ●—●</i>
Mary will be drawing a circle <i>Event, Prog: ●—●</i>	>	<	= oi s d f	John writes a letter <i>Event, Simp: ●</i>
Mary will be drawing a circle <i>Event, Prog: ●—●</i>	"	"	"	John laughs <i>Process, Simp: ●</i>
John will be laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary draws a circle <i>Event, Simp: ●</i>
John will be laughing <i>Process, Prog: ●—●</i>	"	"	"	Mary walks <i>Process, Simp: ●</i>
Mary will be drawing a circle <i>Event, Prog: ●—●</i>	> f oi	< fi o	s d f = o	John is angry <i>State, Simp: ●—●</i>
Mary will be laughing <i>Process, Prog: ●—●</i>	"	"	"	John is angry <i>State, Simp: ●—●</i>
Mary will draw a circle <i>Event, Simp: ●</i>	> f oi	<	= o oi s d si	John is writing a letter <i>Event, Prog: ●—●</i>
Mary will draw a circle <i>Event, Simp: ●</i>	"	"	"	John is laughing <i>Process, Prog: ●—●</i>
John will laugh <i>Process, Simp: ●</i>	"	"	"	Mary is drawing a circle <i>Event, Prog: ●—●</i>
John will laugh <i>Process, Simp: ●</i>	"	"	"	Mary is walking <i>Process, Prog: ●—●</i>
Mary will draw a circle <i>Event, Simp: ●</i>	>	<	= s d o oi si	John writes a letter <i>Event, Simp: ●</i>
Mary will draw a circle <i>Event, Simp: ●</i>	"	"	"	John laughs <i>Process, Simp: ●</i>
John will laugh <i>Process, Simp: ●</i>	"	"	"	Mary draws a circle <i>Event, Simp: ●</i>
John will laugh <i>Process, Simp: ●</i>	"	"	"	Mary walks <i>Process, Simp: ●</i>
Mary will draw a circle <i>Event, Simp: ●</i>	>	<	s d f =	John is angry <i>State, Simp: ●—●</i>
Mary will laugh <i>Process, Simp: ●</i>	"	"	"	John is angry <i>State, Simp: ●—●</i>
John will be angry <i>State, Simp: ●—●</i>	> f oi	< fi o	= oi s d f	Mary is drawing a circle <i>Event, Prog: ●—●</i>
John will be angry <i>State, Simp: ●—●</i>	"	"	"	Mary is walking <i>Process, Prog: ●—●</i>
John will be angry <i>State, Simp: ●—●</i>	>	<	= s d f	Mary draws a circle <i>Event, Simp: ●</i>
John will be angry <i>State, Simp: ●—●</i>	"	"	"	Mary walks <i>Process, Simp: ●</i>
John will be angry <i>State, Simp: ●—●</i>	> f oi	< fi o	= s d f	Mary is happy <i>State, Simp: ●—●</i>

The corresponding selection table for the Future/Present tense combination is given here:

AFTER												
	=	o	oi	s	si	d	di	m	mi	f	fi	<
Dp/Dp			Y							Y		Y
Dp/Ds												Y
Dp/Ss			Y							Y		Y
Ds/Dp			Y							Y		Y
Ds/Ds												Y
Ds/Ss												Y
Ss/Dp			Y							Y		Y
Ss/Ds												Y
Ss/Ss			Y							Y		Y

BEFORE												
	=	o	oi	s	si	d	di	m	mi	f	fi	<
Dp/Dp		Y									Y	Y
Dp/Ds												Y
Dp/Ss		Y									Y	Y
Ds/Dp												Y
Ds/Ds												Y
Ds/Ss												Y
Ss/Dp		Y									Y	Y
Ss/Ds												Y
Ss/Ss		Y									Y	Y

WHILE												
	=	o	oi	s	si	d	di	m	mi	f	fi	<
Dp/Dp	Y		Y	Y		Y				Y		
Dp/Ds	Y	Y		Y		Y				Y		
Dp/Ss	Y		Y	Y		Y				Y		
Ds/Dp	Y	Y	Y	Y	Y	Y				Y		
Ds/Ds	Y	Y	Y	Y	Y	Y				Y		
Ds/Ss	Y			Y		Y				Y		
Ss/Dp	Y		Y	Y		Y				Y		
Ss/Ds	Y			Y		Y				Y		
Ss/Ss	Y			Y		Y				Y		

D Detailed Corpus Analysis

The results of our corpus-based experiment on the LOB indicate that matrix/adjunct pairs that are disallowed by the CDTS occurred rarely in comparison to allowed matrix/adjunct pairs. Thus, our analysis provides strong support for the use of Hornstein’s CDTS in generating surface sentences. However, it is worthwhile to examine the disallowed cases, so that we have a better understanding of where (and why) the current approach would fall short.

Only three cases occurred with higher-than-average frequency: Past/Present (15 times), Present/Past (29 times), and Present Perfect/Past (21 times). In the examples below, we have annotated the sentences with uppercase for matrix/adjunct verbs and brackets for temporal connectives. Sentence numbers are taken from the position of the sentence in the LOB corpus. We have hand-annotated misanalyzed sentences with error comments.

D.1 Past/Present

Of the the Past/Present CDTS violations, 8 were cases where our Perl script was not robust enough to assign the appropriate structure, and the remaining 7 were true violations of the CDTS. All past/present violations are listed at the end of this section.

An example of a case where our Perl script was not robust was the handling of conjunction, as in case 20097: *Their three boys, now successful men, WERE in our children’s Church from the outset, and [when] we DON’t see one another we do not forget.* Other cases of non-robustness were: attachment of temporal clauses to the matrix clause rather than to the nouns they modify (as in case 45809: *... the old car RELAXED like a horse [when] the race IS done*); quotations (as in case 50918: *I THOUGHT [while] you’RE here ...*); and connectives that are used causally rather than temporally (as in case 49509: *... why you TOOK the case, [when] you never TOUCH anything of the sort ...*).

An example of a true CDTS violation is case (28543): *During the latter part of May and early in June the weather was unusually cold and wet, and growth was CHECKED at a time [when] the quality teas of the year ARE made.* Hornstein discusses the use of the present tense to provide a generic or habitual interpretation: it is understood that quality teas are (usually) made at a certain time of the year. In some sense, the adjunct clause is *nontemporal*. (See (Hornstein, 1990, p. 206, fn.20) for related discussion.) In any case, the example above is clearly a case that violates the CDTS, but this type of violation typically coincides with a generic or habitual interpretation. As it turns out, the other 6 violating sentences have a similar nature to them.

The number of true violations (7) is small enough to consider them insignificant. From the point of view of generation, this low count is a clear indicator that we have little to lose by relying on the CDTS for examples such as the one above. In such a case, our generator would select the past tense, thus losing the habitual reading, but still retaining the appropriate temporal ordering. The Past/Present cases are given below:

- **Conjunction ignored:**

(20097) Their three boys, now successful men, WERE in our children’s Church from the outset, and [when] we DON’t see one another we do not forget. ERROR: Conjoined by AND, so separate sentence.
- **Embeddedness (parsing) error:**

(22437) We only SOLD out [when] compelled to DO so by Luigi’s death.
ERROR: to do is a subordinate clause.

(45809) The needle flickered down to sixty, to fifty, and the old car RELAXED like a horse [when] the race IS done. ERROR: when clause is a subordinate inside of the NP headed by horse

(4501) Of the Eisenhower ban announced over the weekend, and six days [before] he LEAVES office, one big dealer SAID: ERROR: before clause is attached as adjunct to VP containing *announced*
- **Quoted expression:**

(50918) I THOUGHT [while] you’RE here. . . .
- **Connective is not truly temporal:**

(37043) The majesty, the familiarity, of these buildings SEEMED to add solemnity to my rite, [as when] old patriarchs COME to grace a marriage. ERROR: *as when* not a temporal connective

(49509) Why you TOOK the case, [when] you never TOUCH anything of the sort. ERROR: *when* used causally (*even though*).

(1053) And [since] the duke IS the landlord of the building (rent 1 s a year) he WAS the obvious choice as guest of honour. ERROR: *since* used causally (*because*).
- **Real violations:**

(231) King Freddie and three other hereditary rulers of native kingdoms in Uganda ARRIVED for talks with colonial Secretary Mr Iain Macleod, [before] the Uganda constitutional conference OPENS next Monday.

(19124) Good morning in verse 1 RESULTED from the closing of the village school, [since when] the children GO to Buckingham and no longer have a holiday on May day.

(4423) Investments and cash at bankers AGGREGATED 3,880,000, representing 16.6 per cent of total assets, [while] reserve funds ARE 4.26 per cent of assets.

(18263) The occasion was, of course, the quatercentenary of the Scottish reformation, but besides this her majesty WAS the very first sovereign lady to honour the Fathers and brethren with her presence, a circumstance not lacking in significance, especially [when] one RECALLS John Knox’s well kent fulminations against women in general and female rulers in particular.

(49625) Beryl’s life recently the whole thing WAS very strange [when] you THINK of it.

(23681) Her life was on the whole unfortunate, and her end sad; yet she WAS a fascinating personality and a fine actress, [while] her life-story IS highly romantic.

(28543) During the latter part of May and early in June the weather was unusually cold and wet, and growth was CHECKED at a time [when] the quality teas of the year ARE made.

D.2 Present/Past

In the case of the CDTS violation involving the Present/Past combination, the same type of robustness issues arise, i.e., attachment errors and lack of handling

for conjunction and quotations. Even so, removing these cases still leaves 12 violations. Among these sentences, *since* is predominately the most frequently used connective. (There is also one case with *before*, one with *when*, and one with *while*; these have a semantics similar in nature to the *since* cases.)

An example of a true CDTS violation is case (3278): *This three-day visit IS President Kennedy's first to Europe [since] he TOOK office*. Hornstein discusses such cases extensively (see, e.g., footnote 13, p. 205), arguing that *since* has a causal, not temporal, meaning in such cases. The truth of this claim is a matter of much debate, even among Hornstein himself (see, for example Hornstein (1977) vs. Hornstein (1990)). But what IS clear is that the connective *since* requires special handling. Our temporal connective tables provide the flexibility for accommodating special cases since they are based primarily on the semantics of the connective rather than on the syntax of the tense structures. However, since our task is to generate, not analyze, sentences, we have the freedom to select a more fortuitous tense combination. Thus, we continue to adopt the approach of using the CDTS to narrow down the choices. A later extension might be to use the connective tables as a fallback mechanism if the CDTS rules out all other options (as would be the case above). The Present/Past cases are given below.

– **Conjunction ignored:**

- (19111) Marsworth also MAKES an interesting reference to the Tring chimney sweeps who come a-dancing all May-day, which refers to the Jack-in-the-Green, the May garland in the far-off days of the little climbing boys and in still further off days [when] the dancer in it REPRESENTED the spirit of vegetation visiting each house to bring fertility in the coming year.

– **Embeddedness (parsing) error:**

- (14422) we UNDERSTAND that, [while] it EXCITED much attention, it did not intrude in any way on the dancing. ERROR: *while clause* in different subordinate relationship
- (16686) I RECALL a matrimonial case of some ten years ago [when] I DID not follow this principle. ERROR: *when clause* attached to *case*
- (4614) [while] dancing in a five past eight revue in Glasgow she WAS called on to DO some swimming in a royal command performance. ERROR: *while clause* does not contain *was*
- (24857) it TAKES its roots in pre-history [when] man, coping with hostile forces, FELT a primal sympathy for his fellow man and sought to relieve his suffering. ERROR: *when clause* attached to *history*
- (27092) but its local government structure, inherited from the days [when] London WAS much smaller, in no way REFLECTS that unity. ERROR: *when clause* attached to *days*
- (1639) it IS Morris and company 1861-1940, a tribute to Morris and his associates 100 years [after] they STARTED their firm. ERROR: *after* attached to *years*
- (18256) the hereditary MacDougall pipers, [while] not so famous as the MacCrimmons of Skye, WERE players and composers of distinction, and the tune, lament for Captain MacDougall, IS one of delicacy and feeling. ERROR: *while* attached to *pipers*
- (29765) this conclusion IS perhaps rather unexpected in view of the appreciable delay in the breakdown of normal electrical activity obtained [when] the insect nervous

- system WAS exposed to solutions of high potassium concentration (Hoyle, 1953 ; Twarog & Roeder, 1956). ERROR: *when* attached to *obtained*
- (282) with the Prime Minister sunning himself in Jamaica and his cabinet out in the grass roots making 160 speeches in 80 constituencies in 10 days, the liberal party ARE holding a national conference here with some 2,000 delegates, the biggest gathering [since] 1958 when Mr Lester Pearson WAS chosen as party leader. ERROR: *since* attached to *gathering*
- (24295) the procedure for private bills is virtually extinct, though there ARE some instances of its use, as in the recent case of the Esso petroleum bill, [when] a private company SOUGHT powers of compulsory purchase. ERROR: *when* attached to *case*
- (24987) gone ARE the days [when] Cossacks GALLOPED across the grassy steppe on superb horses. ERROR: *when* attached to *days*
- (29518) it IS now 50 years [since] Rutherford, working in Manchester, CONCEIVED the idea that the atom had a small concentrated nucleus, and from this idea sprang the whole of our present-day knowledge of atomic structure and our exploitation of its consequences. ERROR: *since* attached to *years*
- (37015) or [when] I SAW her CUT the napkins in two with the Samurai sword? ERROR: *cut* embedded inside of *saw* clause
- (41253) round and round and round, [while] meditatively, as a cow chewing the cud, he LET his eyes REST on the flat water ahead of him. ERROR: *rest* embedded inside of *let* clause
- **Quoted expression:**
- (4846) och, come on, United GROAN [before] Tommy Neilson MADE the vocalists happy by beating Brown.
- (49848) well, dearest one, it IS like this, [when] I LEFT my apartment in London for a short holiday I only drew from my bank enough cash to last me about three weeks.
- **Real violations:**
- (21848) I LOOK back now with great affection on those days of motor-bicycle competition in Edwardian times, [before] I WAS afflicted by the car bug.
- (46990) it 'S only six months [since] we WERE serving together under La Cruz.
- (303) certainly, he IS now a much tougher character politically than [when] he TOOK over the leadership.
- (3278) this three-day visit IS President Kennedy's first to Europe [since] he TOOK office.
- (3722) Cingle, who is under orders for the round tower handicap (3-30), is expected to BECOME Jack Langley's first winner [since] he TOOK charge of Mr W J Weston-Evans' horses at Herringswell Manor.
- (7191) it IS more than two years [since] the society IMPOSED its embargo on the entry of apprentices into the yards because of unemployment among its adult members.
- (19360) it IS some years ago [since] I first BECAME interested in the possible effect of modern noises on animals.
- (30880) the correlations experiment IS too easy for secondary, but not for primary pupils, compared with the other eight experiments ; [while] the projection of shadows test PLACED too many subjects at stage 2B.
- (31382) this IS particularly the case [when] the dead person LIVED to a great age or had high prestige for some reason among his kindred or in the locality.
- (47674) it 'S a long time [since] I FOUGHT a Viet.
- (53274) it 'S a while [since] you SAW me last, the girl reminded her smilingly.
- (46487) you LOOK much more tired [since] you TOOK on that new job.

D.3 Present Perfect/Past

The Present Perfect/Past combination was, by far, the most prevalent CDTS violation. Of the 21 cases, only 3 were a result of non-robust handling by the Perl script. What is interesting, however, is that the interpretation of the violating sentences seems to be similar in nature to that of the Present/Past violations above. Here again, the *since* connective is the most frequently used. In cases where a different connective is used (5 use *when* and 1 uses *while*), the interpretation of the sentences is similar to the *since* cases.

An example of a true CDTS violation is case (6851): *It HAS more than DOUBLED [since] the service STARTED*. The discussion above concerning the distinction between causal vs. temporal interpretations is applicable here. Given that this shows up across different tense pairs, it would appear that using the connective tables as a fallback to accommodate idiosyncracies of the individual connectives (as described above) might be a profitable extension to the current framework. The Perfect/Past cases are given below.

– **Embeddedness (parsing) error:**

- (20095) Mr and Mrs Charlton HAVE BEEN from the first difficult years of war, [when] most lives WERE upset and some tempers were easily frayed, the most loyal and devoted friends. ERROR: *when* attached to *war*
- (13327) Brown HAS HELD his crown since August 1956, [when] he OUTPOINTED Wallace Bud Smith. ERROR: *when* attached to *August*
- (15500) of the men reaching fifty years of age [since] the scheme STARTED, 125 (37.2 per cent) HAVE TAKEN part. ERROR: *since* attached to *age*

– **Real violations:**

- (21827) I HAVE never FORGOTTEN this kindly and thoughtful gesture of Rootes at a time [when] things WERE not going so well for me.
- (2244) he HAS PUBLISHED half a dozen books of poetry and achieved a wider reputation [when] he WROTE the lyrics for the royal court Theatre musical the lily-white boys.
- (2513) I HAVE WARNED the country again and again of this [since] I BECAME Chancellor.
- (6253) [since] price restraint BECAME operative the industry HAS WON success in export markets.
- (6851) it HAS more than DOUBLED [since] the service STARTED.
- (7968) also, nearly a million people HAVE BEEN RE-HOUSED from slums [since] the government's drive STARTED in 1956.
- (13468) there HAS never BEEN a time [when] wines from so many different countries WERE available in Britain.
- (18268) many Commissioners have come from the ranks of the aristocracy and professional classes, some HAVE BEEN personally associated with the work of the kirk, [while] one, James, Duke of Albany and York, brother of Charles 2, WAS a convert to Roman Catholicism.
- (21420) how strange my fate HAS BEEN [since] we WERE together in Brighton with the Regent !
- (28013) [since] television WAS introduced the increase in the power of the stations and the improved sensitivity of receivers HAVE MADE outdoor aerials less necessary in many locations.

- (8866) countless old Dickensian hacks HAVE BEEN bemoaning Pickwick and Micawber ever [since] novelists and critics first BEGAN their resolute march in a different direction.
- (17243) in her long career Miss Horrocks HAS KNOWN only one marriage hitch last summer, [when] ex-assistant-hangman Brian Allen and his Spanish bride Angela Corillo WENT through a marriage ceremony at a Roman catholic church, but forgot to inform Miss Horrocks.
- (29706) outside the Vale of Wardour proper, the Warminster Greensand beds at the base of the Chalk Marl HAVE RECEIVED attention from Jukes-Browne in 1896, 1900-4 and 1901, and from Scanes, jointly with Jukes-Browne, in 1901, and with Pope-Bartlett in 1916 [when], in the latter year, both authors LED the third geologists' association excursion.
- (33887) I HAVE FOLLOWED the progress of this talented young artist's work [since] I CALLED on him, shortly after the war, in his tiny, drab studio in the squalid La Ruche building way over in the 15th arrondissement.
- (36434) and what HAVE you BEEN DOING [since] we MET last time? she asked Erich, more by way of starting a conversation with him than from a desire to know.
- (46077) but the road HAS N'T BEEN USED [since] you LEFT.
- (49260) I 'VE BEEN WORRIED about that damned clip ever [since] I LOST it.
- (28062) [since] March, 1958, when the total employed WAS 542, the council HAVE MADE slight reductions each year bringing about a total decrease of 5 per cent in their staff up to the end of March, 1961.

References

- Allen, James. 1983. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM*, 26(11):832-843.
- Allen, James. 1984. Towards a General Theory of Action and Time. *Artificial Intelligence*, 23(2):123-160.
- Amsili, Pascal and Corinne Rossari. 1998. Tense and Connective Constraints on The Expression of Causality. In *Proceedings of the Thirty-Sixth Annual Meeting of the Association for Computational Linguistics and Seventeenth International Conference on Computational Linguistics*, pages 48-54. Montreal, Canada.
- Androutsopoulos, Ion. 1996. *A Principled Framework for Constructing Natural Language Interfaces To Temporal Databases*. Ph.D. thesis, University of Edinburgh, Edinburgh, UK.
- Androutsopoulos, Ion. 1999. Time, Tense and Aspect in Natural Language Database Interfaces. *Natural Language Engineering*, 4(3):229-276.
- Bach, E. and R. Harms. 1968. *Universals in Linguistic Theory*. Holt, Rinehart and Winston, New York.
- Baldwin, Jennifer A. 2002. Learning Temporal Annotation of French News. Master's thesis, Graduate School of Arts and Sciences, Georgetown University, Washington, DC.
- Bennett, W., T. Herlick, K. Hoyt, J. Liro, and A. Santisteban. 1990. A Computational Model of Aspect and Verb Semantics. *Machine Translation*, 4(4):247-280.

- Brent, Michael. 1990. A Simplified Theory of Tense Representations and Constraints on Their Composition. In *Proceedings of the 28th Annual Conference of the Association for Computational Linguistics*, University of Pittsburgh, Pittsburgh, PA.
- Comrie, Bernard. 1976. *Aspect*. Cambridge University Press, Cambridge, MA.
- Crouch, R. and S. Pulman. 1993. Time and modality in a natural language interface to a planning system. *Artificial Intelligence*, 63:265–304.
- Dale, Robert. 1992. *Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes*. MIT Press, Cambridge, MA.
- Danlos, Laurence. 1999. Event Coreference Between Two Sentences. In *Proceedings of the Third International Workshop on Computational Semantics (IWCS'99)*, Tilburg.
- Danlos, Laurence. 2000. Event Coreference in Causal Discourses. In *Meaning of Words*. Cambridge University Press, Cambridge, England.
- Dorr, Bonnie J. 1992. A Parameterized Approach to Integrating Aspect with Lexical-Semantics for Machine Translation. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, pages 257–264, University of Delaware, Newark, DE.
- Dorr, Bonnie J. 2001. LCS Verb Database. Technical Report Online Software Database, University of Maryland, College Park, MD. http://www.umiacs.umd.edu/~bonnie/LCS_Database_Documentation.html.
- Dorr, Bonnie J. and Terry Gaasterland. 1995. Selecting Tense, Aspect, and Connecting Words in Language Generation. In *Proceedings of IJCAI-95*, pages 1299–1305, Montreal, Canada.
- Dorr, Bonnie J., Nizar Habash, and David Traum. 1998. A Thematic Hierarchy for Efficient Generation from Lexical-Conceptual Structure. In *Proceedings of the Third Conference of the Association for Machine Translation in the Americas, AMTA-98, in Lecture Notes in Artificial Intelligence, 1529*, pages 333–343, Langhorne, PA, October 28–31.
- Dorr, Bonnie J. and Mari Broman Olsen. 1996. Multilingual Generation: The Role of Telicity in Lexical Choice and Syntactic Realization. *Machine Translation*, 11(1–3):37–74.
- Dorr, Bonnie J. and Clare R. Voss. 1996. A Multi-Level Approach to Interlingual MT: Defining the Interface Between Representational Languages. *International Journal of Expert Systems*, 9(1):15–51.
- Dowty, David. 1979. *Word Meaning in Montague Grammar*. Reidel, Dordrecht.
- Elhadad, Michael and Kathy McKeown. 1990. Generating Connectives. In *Proceedings of the 13th International Conference on Computational Linguistics (COLING 1990)*, pages 97–101, Helsinki, Finland.
- Ferro, Lisa, Inderjeet Mani, Beth Sundheim, and George Wilson. 2000. TIDES Temporal Annotation Guidelines Draft Version 1.0. MITRE Technical Report MTR 00W0000094, The MITRE Corporation, McLean, Virginia.
- Filatova, Elena and Eduard Hovy. 2001. Assigning Time-Stamps To Event-Clauses. In *Proceedings of ACL'2001 Workshop on Temporal and Spatial Information Processing*, Toulouse, France.

- Gagnon, Michel and Guy Lapalme. 1996. From Conceptual Time to Linguistic Time. *Computational Linguistics*, 22(1):91–127.
- Galton, A. 1990. A Critical Examination of Allen’s Theory of Action and Time. *Artificial Intelligence*, 42:159–188.
- Habash, Nizar. 2002. Generation-Heavy Machine Translation. In *Proceedings of the International Natural Language Generation Conference (INLG’02) Student Session*, New York.
- Habash, Nizar and Bonnie J. Dorr. 2002. Handling Translation Divergences: Combining Statistical and Symbolic Techniques in Generation-Heavy Machine Translation. In *Fifth Conference of the Association for Machine Translation in the Americas, AMTA-2002*, Tiburon, California.
- Hornstein, Norbert. 1990. *As Time Goes By*. The MIT Press, Cambridge, MA.
- Hovy, Eduard H. 1993. Automated Discourse Generation Using Discourse Structure Relations. *Artificial Intelligence*, 63(1&2):341–384.
- Hwang, Chung Hee and Lenhart Schubert. 1994. Interpreting Tense, Aspect and Time Adverbials: A Compositional, Unified Approach. In *Proceedings of the First International Conference, ICTL’94 on Temporal Logic*, pages 238–264, Bonn, Germany, July.
- Jackendoff, Ray. 1983. *Semantics and Cognition*. The MIT Press, Cambridge, MA.
- Jackendoff, Ray. 1990. *Semantic Structures*. The MIT Press, Cambridge, MA.
- Kasper, Robert T. 1989. A Flexible Interface for Linking Applications to PENMAN’s Sentence Generator. In *Proceedings of the DARPA Workshop on Speech and Natural Language*. Available from USC/Information Sciences Institute, Marina del Rey, CA.
- Lascarides, Alex and Jon Oberlander. 1993. Temporal Connectives in a Discourse Context. In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL93)*, pages 260–268, Utrecht, The Netherlands.
- Lesperance, Y. and H. Levesque. 1990. Indexical Knowledge in Robot Plans. In *Proceedings 8th National Conference on Artificial Intelligence, AAAI-90*.
- Mani, Inderjeet and George Wilson. 2000. Robust Temporal Processing of News. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL’2000)*, pages 69–76, Hong Kong.
- Merriam-Webster. 1976. *Webster’s Seventh New Collegiate Dictionary*. on line version, G. & C. Merriam Co., Springfield, MA.
- Moens, Marc and Mark Steedman. 1988. Temporal Ontology and Temporal Reference. *Computational Linguistics: Special Issue on Tense and Aspect*, 14(2):15–28.
- Mourelatos, A. 1981. Events, Processes, and States. In P. Tedeschi and A. Zaenen, editors, *Syntax and Semantics 14: Tense and Aspect*. Academic Press.
- Nirenburg, Sergei and James Pustejovsky. 1988. Processing Aspectual Semantics. In *Proceedings of Tenth Annual Conference of the Cognitive Science Society*, pages 658–665, Montreal, Canada.
- Olsen, Mari Broman. 1994. *A Semantic and Pragmatic Model of Lexical and Grammatical Aspect*. Ph.D. thesis, Northwestern University, Evanston, IL.

- Olsen, Mari Broman. 1997. *A Semantic and Pragmatic Model of Lexical and Grammatical Aspect*. Garland, New York.
- Olsen, Mari Broman, David Traum, Carol van Ess-Dykema, and Amy Weinberg. 2000. Telicity as a Cue to Temporal and Discourse Structure in Chinese-English Machine Translation. In *Interlingual Workshop, Proceedings of the 38th Annual Conference of the Association for Computational Linguistics (ACL)*, Seattle, Washington.
- Olsen, Mari Broman, David Traum, Carol van Ess-Dykema, and Amy Weinberg. 2001. Implicit Cues for Explicit Generation: Using Telicity as a Cue for Tense Structure in a Chinese to English MT System. In *Proceedings of MT Summit VIII, Santiago de Compostella, Spain*.
- Passoneau, Rebecca. 1988. A Computational Model of the Semantics of Tense and Aspect. *Computational Linguistics: Special Issue on Tense and Aspect*, 14(2):44–60.
- Pustejovsky, J., S. Bergler, and P. Anick. 1993. Lexical Semantic Techniques for Corpus Analysis. *Computational Linguistics*, 19(2).
- Pustejovsky, James. 1988. The Geometry of Events. Technical report, Massachusetts Institute of Technology, Center for Cognitive Science, Cambridge, MA.
- Pustejovsky, James. 1991a. The Generative Lexicon. *Computational Linguistics*, 17(4):409–441.
- Pustejovsky, James. 1991b. The Syntax of Event Structure. In Beth Levin and Steven Pinker, editors, *Lexical and Conceptual Semantics*. Blackwell, Cambridge, MA, pages 47–81.
- Pustejovsky, James, Roser Saurí, Andrea Setzer, Robert Gaizauskas, and Bob Ingria. 2002. TimeML Annotation Guidelines. In *Proceedings of ARDA Workshop on Advanced Question Answering Technology (TERQAS: Time and Event Recognition for Question Answering)*, Brandeis University, Waltham, MA. <http://www.cs.brandeis.edu/~jamesp/arda/time/documentation/AnnotationGuideline-v0.4.0.pdf>.
- Pustejovsky, James, Janyce Wiebe, and Mark Maybury. 2002. Multiple Perspective and Temporal Question Answering. In *Proceedings of LREC Workshop on Question Answering: Strategy and Resources*, pages 39–43, Las Palmas, Canary Islands, Spain.
- Radev, Dragomir and Beth Sundheim. 2002. Using TimeML in Question Answering. Technical report, Brandeis University, Waltham, MA.
- Reichenbach, H. 1947. *Elements of Symbolic Logic*. Macmillan, London.
- Reiter, Ehud and Robert Dale. 1997. Building Applied Natural Language Generation Systems. *Journal of Natural Language Engineering*, 3(1):57–87.
- Schilder, Frank. 1999. Presupposition Triggered by Temporal Connectives. In *Proceedings of Atelier Thematique TALN, Institut d’Etudes Scientifiques, Cargese, Corse*.
- Schilder, Frank and Christopher Habel. 2001. From Temporal Expressions to Temporal Information: Semantic Tagging of News Messages. In *Proceedings of ACL’2001 Workshop on Temporal and Spatial Information Processing*, Toulouse, France.

- Setzer, Andrea. 2001. *Temporal Information in Newswire Articles: An Annotation Scheme and Corpus Study*. Ph.D. thesis, University of Sheffield, Sheffield, UK.
- Setzer, Andrea. 2002. On the Importance of Annotating Temporal Event-Event Relations in Text. In *Proceedings of LREC Workshop on Annotation Standards for Temporal Information in Natural Language*, pages 52–60, Las Palmas, Canary Islands, Spain.
- Setzer, Andrea and Robert Gaizauskas. 2001. Pilot Study On Annotating Temporal Relations In Text. In *Proceedings of ACL'2001 Workshop on Temporal and Spatial Information Processing*, Toulouse, France.
- Snodgrass, Richard T. 1990. Research Concerning Time in Databases: Project Summaries. *ACM SIGMOD Record*, 15(4):19–39.
- Snodgrass, Richard T. 1995. *The TSQL2 Temporal Query Language*. Kluwer, Norwell, MA.
- Steedman, Mark. 1997. Temporality. In *Handbook of Logic and Language*. Elsevier Science, North-Holland, pages 179–239.
- Torp, Kristian, Christian S. Jensen, and Richard T. Snodgrass. 2000. Effective Timestamping in Databases. *VLDB Journal: Very Large Data Bases*, 8(3–4):267–288.
- Vendler, Zeno. 1967. Verbs and Times. *Linguistics in Philosophy*, pages 97–121.
- Verkuyl, Henk. 1972. *On the Compositional Nature of the Aspects*. *Foundations of Language Supplementary Series 15*. Reidel, Dordrecht.
- Vilain, M., H. Kautz, and P. van Beek. 1990. Constraint Propagation Algorithms for Temporal Reasoning: A Revised Report. In *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann, San Mateo, CA.
- Williams, Brian. 1990. Doing Time: Putting Qualitative Reasoning on Firmer Ground. In *Readings in Qualitative Reasoning about Physical Systems*, Morgan Kaufmann, San Mateo, CA.
- Wilson, George, Inderjeet Mani, Beth Sundheim, and Lisa Ferro. 2001. A Multilingual Approach to Annotating and Extracting Temporal Information. In *Proceedings of the ACL'2001 Workshop on Temporal and Spatial Information Processing*, Toulouse, France.
- Yip, Kenneth. 1985. Tense, Aspect and the Cognitive Representation of Time. In *Proceedings of the 23rd Annual Conference of the Association for Computational Linguistics*, pages 18–26, Chicago, IL.